

PATENT COOPERATION TREATY

PCT

NOTIFICATION OF ELECTION

(PCT Rule 61.2)

From the INTERNATIONAL BUREAU

To:

Commissioner
US Department of Commerce
United States Patent and Trademark
Office, PCT
2011 South Clark Place Room
CP2/5C24
Arlington, VA 22202
ETATS-UNIS D'AMERIQUE
in its capacity as elected Office

Date of mailing (day/month/year) 06 July 2001 (06.07.01)	
International application No. PCT/US00/24000	Applicant's or agent's file reference 10559-308WO1
International filing date (day/month/year) 31 August 2000 (31.08.00)	Priority date (day/month/year) 01 September 1999 (01.09.99)
Applicant WOLRICH, Gilbert et al	

1. The designated Office is hereby notified of its election made:

☒ in the demand filed with the International Preliminary Examining Authority on:

30 March 2001 (30.03.01)

☐ in a notice effecting later election filed with the International Bureau on:

2. The election ☒ was
☐ was not

made before the expiration of 19 months from the priority date or, where Rule 32 applies, within the time limit under Rule 32.2(b).

The International Bureau of WIPO 34, chemin des Colombettes 1211 Geneva 20, Switzerland	Authorized officer Nestor Santesso
Facsimile No.: (41-22) 740.14.35	Telephone No.: (41-22) 338.83.38

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
8 March 2001 (08.03.2001)

PCT

(10) International Publication Number
WO 01/16714 A1

(51) International Patent Classification⁷: **G06F 9/312,**
9/315, 9/38

(71) Applicant (for all designated States except US): **INTEL CORPORATION** [US/US]; 2200 Mission College Boulevard, Santa Clara, CA 95052 (US).

(21) International Application Number: **PCT/US00/24000**

(22) International Filing Date: **31 August 2000 (31.08.2000)**

(25) Filing Language: **English**

(26) Publication Language: **English**

(30) Priority Data:
60/151,961 1 September 1999 (01.09.1999) **US**

(63) Related by continuation (CON) or continuation-in-part (CIP) to earlier application:
US 60/151,961 (CIP)
Filed on 1 September 1999 (01.09.1999)

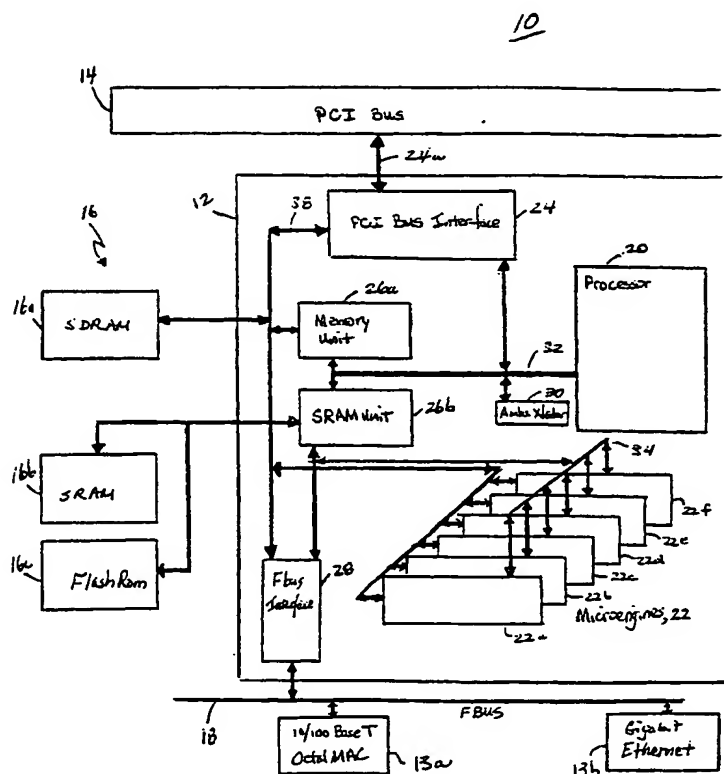
(72) Inventors; and
(75) Inventors/Applicants (for US only): **WOLRICH, Gilbert** [US/US]; 4 Cider Mill Road, Framingham, MA 01701 (US). **ADILETTA, Matthew, J.** [US/US]; 20 Monticello Drive, Worcester, MA 01603 (US). **WHEELER, William** [US/US]; 745 School Street, Webster, MA 01570 (US). **BERNSTEIN, Debra** [US/US]; 38 Helen Street, Waltham, MA 02452 (US). **HOOPER, Donald** [US/US]; 19 Main Circle, Shrewsbury, MA 01545 (US).

(74) Agents: **MALONEY, Denis, G.**; Fish & Richardson P.C., 225 Franklin Street, Boston, MA 02110-2804 et al. (US).

(81) Designated States (national): **AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CR, CU, CZ,**

[Continued on next page]

(54) Title: **FAST WRITE INSTRUCTION FOR MICRO ENGINE USED IN MULTITHREADED PARALLEL PROCESSOR ARCHITECTURE**



(57) Abstract: A method of operating a processor (12) including receiving data in a processing thread having a processing thread number and shifting the data into a register corresponding to the processing thread number.

WO 01/16714 A1



DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW.

- (84) **Designated States (regional):** ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

Published:

- *With international search report.*
- *Before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments.*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

FAST WRITE INSTRUCTION FOR MICRO ENGINE USED
IN MULTITHREADED PARALLEL PROCESSOR ARCHITECTURE

5

BACKGROUND

This invention relates to a memory instruction for computer processors.

Parallel processing is an efficient form of information processing of concurrent events in a computing process. Parallel processing demands concurrent execution of many programs in a computer, in contrast to sequential processing. In the context of a parallel processor, parallelism involves doing more than one thing at the same time. Unlike a serial paradigm where all tasks are performed sequentially at a single station or a pipelined machine where tasks are performed at specialized stations, with parallel processing, a number of stations are provided with each capable of performing all tasks. That is, in general all or a number of the stations work simultaneously and independently on the same or common elements of a problem. Certain problems are suitable for solution by applying parallel processing.

DESCRIPTION OF DRAWINGS

The foregoing features and other aspects of the invention will be described further in detail by the accompanying drawings, in which:

FIG. 1 is a block diagram of a communication system employing a hardware-based multithreaded processor.

FIG. 2 is a detailed block diagram of the hardware-based multithreaded processor of FIG. 1.

FIG. 3 is a block diagram of a micro engine functional unit employed in the hardware-based multithreaded processor of FIGS. 1 and 2.

FIG. 4 is a block diagram of a pipeline in the micro engine of FIG. 3.

FIG. 5 is a block diagram illustrating a format for arithmetic logic unit instruction results.

Like reference symbols in the various drawings indicate like elements.

DETAILED DESCRIPTION

Referring to FIG. 1, a communication system 10 includes a parallel, hardware-based multithreaded processor 12. The hardware-based multithreaded processor 12 is coupled to a bus such as a PCI bus 14, a memory system 16 and a second bus 18. The system 10 is especially useful for tasks that can be broken into parallel subtasks or functions. Specifically, hardware-based multithreaded processor 12 is useful for tasks that are bandwidth oriented rather than latency oriented. The hardware-based multithreaded processor 12 has multiple micro engines 22 each with multiple hardware controlled threads that can be simultaneously active and independently work on a task.

The hardware-based multithreaded processor 12 also includes a central controller 20 that assists in loading micro code control for other resources of the hardware-based multithreaded processor 12 and performs other general purpose computer type functions such as handling protocols, exceptions, extra support for packet processing where the micro engines 22 pass the packets off for more detailed processing such as in boundary conditions. In one embodiment, the processor 20 is a Strong Arm® (Arm is a trademark of ARM Limited, United Kingdom) based architecture. The general-purpose microprocessor 20 has an operating system. Through the operating system the processor 20 can call functions to operate on micro engines 22a-22f. The processor 20 can use any supported operating system, preferably a real time operating system. For the core processor 20 implemented as Strong Arm architecture, operating systems such as, Microsoft-NT real-time, VXWorks and μ CUS, a freeware operating system available over the Internet, can be used.

Functional micro engines (micro engines) 22a-22f each maintain program counters in hardware and states associated with the program counters. Effectively, a corresponding number of sets of threads can be simultaneously active on each of the micro engines 22a-22f while only one is actually operating at any one time.

In an embodiment, there are six micro engines 22a-22f as shown. Each micro engine 22a-22f has capabilities for processing four hardware threads. The six micro engines 22a-22f operate with shared resources including memory system 16 and bus interfaces 24 and 28. The memory system 16 includes a Synchronous Dynamic Random Access Memory (SDRAM) controller 26a and a Static Random Access Memory (SRAM) controller 26b. SDRAM memory 16a and SDRAM controller 26a are typically used for processing large volumes of data, e.g., processing of network payloads from network

packets. The SRAM controller 26b and SRAM memory 16b are used in a networking implementation for low latency, fast access tasks, e.g., accessing look-up tables, memory for the core processor 20, and so forth.

1 The six micro engines 22a-22f access either the SDRAM 16a or SRAM 16b based
5 on characteristics of the data. Thus, low latency, low bandwidth data is stored in and
fetched from SRAM 16b, whereas higher bandwidth data for which latency is not as
important, is stored in and fetched from SDRAM 16a. The micro engines 22a-22f can
execute memory reference instructions to either the SDRAM controller 26a or SRAM
controller 16b.

10 Advantages of hardware multithreading can be explained by SRAM or SDRAM
memory accesses. As an example, an SRAM access requested by a Thread_0, from a
micro engine will cause the SRAM controller 26b to initiate an access to the SRAM
memory 16b. The SRAM controller 26b controls arbitration for the SRAM bus, accesses
15 the SRAM 16b, fetches the data from the SRAM 16b, and returns data to a requesting
micro engine 22a-22f. During an SRAM access, if the micro engine, e.g., micro engine
22a, had only a single thread that could operate, that micro engine would be dormant until
data was returned from the SRAM 16b. By employing hardware context swapping within
each of the micro engines 22a-22f, the hardware context swapping enables other contexts
with unique program counters to execute in that same micro engine. Thus, another
20 thread, e.g., Thread_1 can function while the first thread, i.e., Thread_0, is awaiting the
read data to return. During execution, Thread_1 may access the SDRAM memory 16a.
While Thread_1 operates on the SDRAM unit 16a, and Thread_0 is operating on the
SRAM unit 16b, a new thread, e.g., Thread_2 can now operate in the micro engine 22a.
Thread_2 can operate for a certain amount of time until it needs to access memory or
25 perform some other long latency operation, such as making an access to a bus interface.
Therefore, simultaneously, the processor 12 can have a bus operation, SRAM operation
and SDRAM operation all being completed or operated upon by one micro engine 22a
and have one more thread available to process more work in the data path.

30 The hardware context swapping also synchronizes completion of tasks. For
example, two threads could hit the same shared resource e.g., SRAM 16b. Each one of
these separate functional units, e.g., the FBUS interface 28, the SRAM controller 26a, and
the SDRAM controller 26b, when they complete a requested task from one of the micro
engine thread contexts reports back a flag signaling completion of an operation. When

the micro engine receives the flag, the micro engine can determine which thread to turn on.

An application for the hardware-based multithreaded processor 12 is as a network processor. As a network processor, the hardware-based multithreaded processor 12
5 interfaces to network devices such as a media access controller device e.g., a 10/100BaseT Octal MAC 13a or a Gigabit Ethernet device 13b. In general, as a network processor, the hardware-based multithreaded processor 12 can interface to any type of communication device or interface that receives/sends large amounts of data.

Communication system 10 functioning in a networking application could receive network
10 packets from the devices 13a, 13b and process those packets in a parallel manner. With the hardware-based multithreaded processor 12, each network packet can be independently processed.

Another example for use of processor 12 is a print engine for a postscript processor or as a processor for a storage subsystem, e.g., Redundant Array of Independent
15 Disk (RAID) storage, a category of disk drives that employs two or more drives in combination for fault tolerance and performance. A further use is as a matching engine. In the securities industry for example, the advent of electronic trading requires the use of electronic matching engines to match orders between buyers and sellers. These and other parallel types of tasks can be accomplished utilizing the system 10.

20 The processor 12 includes the bus interface 28 that couples the processor to the second bus 18. In an embodiment, bus interface 28 couples the processor 12 to the FBUS (FIFO bus) 18. The FBUS interface 28 is responsible for controlling and interfacing the processor 12 to the FBUS 18. The FBUS 18 is a 64-bit wide FIFO bus, used to interface to Media Access Controller (MAC) devices, e.g., 10/100 Base T Octal MAC 13a.

25 The processor 12 includes a second interface e.g., PCI bus interface 24, that couples other system components that reside on the PCI 14 bus to the processor 12. The PCI bus interface 24 provides a high-speed data path 24a to memory 16, e.g., SDRAM memory 16a. Through PCI bus interface 24 data can be moved quickly from the SDRAM 16a through the PCI bus 14, via direct memory access (DMA) transfers. The hardware
30 based multithreaded processor 12 supports image transfers. The hardware based multithreaded processor 12 can employ DMA channels so if one target of a DMA transfer is busy, another one of the DMA channels can take over the PCI bus 14 to deliver information to another target to maintain high processor 12 efficiency. Additionally, the

PCI bus interface 24 supports target and master operations. Target operations are operations where slave devices on bus 14 access SDRAMs through reads and writes that are serviced as a slave to a target operation. In master operations, the processor core 20 sends data directly to or receives data directly from the PCI interface 24.

5 Each of the functional units 22 is coupled to one or more internal buses. As described below, the internal buses are dual, 32 bit buses (i.e., one bus for read and one for write). The hardware-based multithreaded processor 12 also is constructed such that the sum of the bandwidths of the internal buses in the processor 12 exceed the bandwidth of external buses coupled to the processor 12. The processor 12 includes an internal core
10 processor bus 32, e.g., an ASB Advanced System Bus (ASB), that couples the processor core 20 to the memory controller 26a, 26b and to an ASB translator 30, described below. The ASB bus 32 is a subset of the so-called Advanced Microcontroller Bus Architecture (AMBA) bus that is used with the Strong Arm processor core 20. AMBA is an open standard, on-chip bus specification that details a strategy for the interconnection and
15 management of functional blocks that makes up a System-on-chip (SoC). The processor 12 also includes a private bus 34 that couples the micro engine units 22 to SRAM controller 26b, ASB translator 30 and FBUS interface 28. A memory bus 38 couples the memory controller 26a, 26b to the bus interfaces 24 and 28 and memory system 16 including flashrom 16c that is used for boot operations and so forth.

20 Referring to FIG. 2, each of the micro engines 22a-22f includes an arbiter that examines flags to determine the available threads to be operated upon. Any thread from any of the micro engines 22a-22f can access the SDRAM controller 26a, SDRAM controller 26b or FBUS interface 28. The memory controllers 26a and 26b each include queues to store outstanding memory reference requests. The queues either maintain order
25 of memory references or arrange memory references to optimize memory bandwidth. For example, if a thread_0 has no dependencies or relationship to a thread_1, there is no reason that thread_1 and thread_0 cannot complete their memory references to the SRAM unit 16b out of order. The micro engines 22a-22f issue memory reference requests to the memory controllers 26a and 26b. The micro engines 22a-22f flood the memory
30 subsystems 26a and 26b with enough memory reference operations such that the memory subsystems 26a and 26b become the bottleneck for processor 12 operation.

If the memory subsystem 16 is flooded with memory requests that are independent in nature, the processor 12 can perform memory reference sorting. Memory reference

sorting improves achievable memory bandwidth. Memory reference sorting, as described below, reduces dead time or a bubble that occurs with accesses to SRAM 16b. With memory references to SRAM 16b, switching current direction on signal lines between reads and writes produces a bubble or a dead time waiting for current to settle on
 5 conductors coupling the SRAM 16b to the SRAM controller 26b.

That is, the drivers that drive current on the bus need to settle out prior to changing states. Thus, repetitive cycles of a read followed by a write can degrade peak bandwidth. Memory reference sorting allows the processor 12 to organize references to memory such that long strings of reads can be followed by long strings of writes. This
 10 can be used to minimize dead time in the pipeline to effectively achieve closer to maximum available bandwidth. Reference sorting helps maintain parallel hardware context threads. On the SDRAM 16a, reference sorting allows hiding of pre-charges from one bank to another bank. Specifically, if the memory system 16b is organized into an odd bank and an even bank, while the processor is operating on the odd bank, the memory
 15 controller can start pre-charging the even bank. Pre-charging is possible if memory references alternate between odd and even banks. By ordering memory references to alternate accesses to opposite banks, the processor 12 improves SDRAM bandwidth. Additionally, other optimizations can be used. For example, merging optimizations where operations that can be merged, are merged prior to memory access, open page
 20 optimizations where by examining addresses an opened page of memory is not reopened, chaining, as will be described below, and refreshing mechanisms, can be employed.

The FBUS interface 28 supports Transmit and Receive flags for each port that a MAC device supports, along with an Interrupt flag indicating when service is warranted. The FBUS interface 28 also includes a controller 28a that performs header processing of
 25 incoming packets from the FBUS 18. The controller 28a extracts the packet headers and performs a micro programmable source/destination/protocol hashed lookup (used for address smoothing) in SRAM 16b. If the hash does not successfully resolve, the packet header is sent to the processor core 20 for additional processing. The FBUS interface 28 supports the following internal data transactions:

30

FBUS unit	(Shared bus SRAM)	to/from micro engine.
FBUS unit	(via private bus)	writes from SDRAM Unit.
FBUS unit	(via Mbus)	Reads to SDRAM.

The FBUS 18 is a standard industry bus and includes a data bus, e.g., 64 bits wide and sideband control for address and read/write control. The FBUS interface 28 provides the ability to input large amounts of data using a series of input and output FIFOs 29a-29b. From the FIFOs 29a-29b, the micro engines 22a-22f fetch data from or command the SDRAM controller 26a to move data from a receive FIFO in which data has come from a device on bus 18, into the FBUS interface 28. The data can be sent through memory controller 26a to SDRAM memory 16a, via a direct memory access. Similarly, the micro engines can move data from the SDRAM 26a to interface 28, out to FBUS 18, via the FBUS interface 28.

Data functions are distributed amongst the micro engines 22. Connectivity to the SRAM 26a, SDRAM 26b and FBUS 28 is via command requests. A command request can be a memory request or a FBUS request. For example, a command request can move data from a register located in a micro engine 22a to a shared resource, e.g., an SDRAM location, SRAM location, flash memory or some MAC address. The commands are sent out to each of the functional units and the shared resources. However, the shared resources do not need to maintain local buffering of the data. Rather, the shared resources access distributed data located inside of the micro engines 22a-22f. This enables micro engines 22a-22f, to have local access to data rather than arbitrating for access on a bus and risk contention for the bus. With this feature, there is a zero cycle stall for waiting for data internal to the micro engines 22a-22f.

The data buses, e.g., ASB bus 30, SRAM bus 34 and SDRAM bus 38 coupling these shared resources, e.g., memory controllers 26a and 26b, are of sufficient bandwidth such that there are no internal bottlenecks. In order to avoid bottlenecks, the processor 12 has an bandwidth requirement where each of the functional units is provided with at least twice the maximum bandwidth of the internal buses. As an example, the SDRAM 16a can run a 64 bit wide bus at 83 MHz. The SRAM data bus could have separate read and write buses, e.g., could be a read bus of 32 bits wide running at 166 MHz and a write bus of 32 bits wide at 166 MHz. That is, in essence, 64 bits running at 166 MHz which is effectively twice the bandwidth of the SDRAM.

The core processor 20 also can access the shared resources. The core processor 20 has a direct communication to the SDRAM controller 26a to the bus interface 24 and to SRAM controller 26b via bus 32. However, to access the micro engines 22a-22f and

transfer registers located at any of the micro engines 22a-22f, the core processor 20 access the micro engines 22a-22f via the ASB Translator 30 over bus 34. The ASB translator 30 can physically reside in the FBUS interface 28, but logically is distinct. The ASB Translator 30 performs an address translation between FBUS micro engine transfer register locations and core processor addresses (i.e., ASB bus) so that the core processor 20 can access registers belonging to the micro engines 22a-22f.

Although micro engines 22a-22f can use the register set to exchange data as described below, a scratchpad memory 27 is also provided to permit micro engines 22a-22f to write data out to the memory for other micro engines to read. The scratchpad 27 is coupled to bus 34.

The processor core 20 includes a RISC core 50 implemented in a five stage pipeline performing a single cycle shift of one operand or two operands in a single cycle, provides multiplication support and 32 bit barrel shift support. This RISC core 50 is a standard Strong Arm® architecture but it is implemented with a five-stage pipeline for performance reasons. The processor core 20 also includes a 16-kilobyte instruction cache 52, an 8-kilobyte data cache 54 and a prefetch stream buffer 56. The core processor 20 performs arithmetic operations in parallel with memory writes and instruction fetches. The core processor 20 interfaces with other functional units via the ARM defined ASB bus. The ASB bus is a 32-bit bi-directional bus 32.

Referring to FIG. 3, an exemplary one of the micro engines 22a-22f, e.g., micro engine 22f, is shown. The micro engine 22f includes a control store 70, which, in one implementation, includes a RAM of here 1,024 words of 32 bit. The RAM stores a micro program (not shown). The micro program is loadable by the core processor 20. The micro engine 22f also includes controller logic 72. The controller logic 72 includes an instruction decoder 73 and program counter (PC) units 72a-72d. The four micro program counters 72a-72d are maintained in hardware. The micro engine 22f also includes context event switching logic 74. Context event logic 74 receives messages (e.g., SEQ_#_EVENT_RESPONSE; FBI_EVENT_RESPONSE; SRAM _EVENT_RESPONSE; SDRAM _EVENT_RESPONSE; and ASB _EVENT_RESPONSE) from each one of the shared resources, e.g., SRAM 26a, SDRAM 26b, or processor core 20, control and status registers, and so forth. These messages provide information on whether a requested function has completed. Based on whether or not a function requested by a thread has completed and signaled completion, the thread

needs to wait for that completion signal, and if the thread is enabled to operate, then the thread is placed on an available thread list (not shown). The micro engine 22f can have a maximum of four threads available.

In addition to event signals that are local to an executing thread, the micro engines 22a-22f employ signaling states that are global. With signaling states, an executing thread can broadcast a signal state to all micro engines 22a-22f, e.g., Receive Request Available (RRA) signal, any and all threads in the micro engines 22a-22f can branch on these signaling states. These signaling states can be used to determine availability of a resource or whether a resource is due for servicing.

The context event logic 74 has arbitration for the four threads. In an embodiment, the arbitration is a round robin mechanism. Other techniques could be used including priority queuing or weighted fair queuing. The micro engine 22f also includes an execution box (EBOX) data path 76 that includes an arithmetic logic unit (ALU) 76a and general-purpose register set 76b. The ALU 76a performs arithmetic and logical functions as well as shift functions. The register set 76b has a relatively large number of general-purpose registers. In an embodiment, there are 64 general-purpose registers in a first bank, Bank A and 64 in a second bank, Bank B. The general-purpose registers are windowed so that they are relatively and absolutely addressable.

The micro engine 22f also includes a write transfer register stack 78 and a read transfer stack 80. These registers 78 and 80 are also windowed so that they are relatively and absolutely addressable. Write transfer register stack 78 is where write data to a resource is located. Similarly, read register stack 80 is for return data from a shared resource. Subsequent to or concurrent with data arrival, an event signal from the respective shared resource e.g., the SRAM controller 26a, SDRAM controller 26b or core processor 20 will be provided to context event arbiter 74, which will then alert the thread that the data is available or has been sent. Both transfer register banks 78 and 80 are connected to the execution box (EBOX) 76 through a data path. In an embodiment, the read transfer register has 64 registers and the write transfer register has 64 registers.

Referring to FIG. 4, the micro engine data path maintains a 5-stage micro-pipeline 82. This pipeline includes lookup of microinstruction words 82a, formation of the register file addresses 82b, read of operands from register file 82c, ALU shift or compare operations 82d, and write-back of results to registers 82e. By providing a write-back data bypass into the ALU/shifter units, and by assuming the registers are implemented as a

register file (rather than a RAM), the micro engine 22f can perform a simultaneous register file read and write, which completely hides the write operation.

The SDRAM interface 26a provides a signal back to the requesting micro engine on reads that indicates whether a parity error occurred on the read request. The micro engine micro code is responsible for checking the SDRAM 16a read Parity flag when the micro engine uses any return data. Upon checking the flag, if it was set, the act of branching on it clears it. The Parity flag is only sent when the SDRAM 16a is enabled for checking, and the SDRAM 16a is parity protected. The micro engines 22 and the PCI Unit 14 are the only requestors notified of parity errors. Therefore, if the processor core 20 or FIFO 18 requires parity protection, a micro engine assists in the request. The micro engines 22a-22f support conditional branches.

Referring to FIG. 5, a format for arithmetic logic unit (ALU) instruction is shown. The micro engines 22 support an instruction set. The instruction set includes instructions that perform an operation on one or two operands and deposit the result into the destination register, and update ALU condition codes according to the result of the operation. Condition codes are lost during context swaps.

The computer instruction architecture set also includes a fast write (FAST_WR) instruction. The FAST_WR instruction for immediate data writes specified immediate data to a specified control and status register (CSR) by having write data specified directly in the instruction rather than in a transfer register. FAST_WR writes the specified immediate data to the specified FBI CSR. This improves performance by eliminating the need for the FBI unit to pull the data from a transfer register. The FBI unit automatically shifts the immediate data into the appropriate register field corresponding to the thread that is writing the FAST_WR data.

A format for the fast write instruction is: fast_wr [immed_data, csr_addr], optional_token, where a description of each of the fields follows.

The "immed_data" field represents 10 bits of the immediate data to be written to the control and status register (CSR); valid immed_data values are 0 through 0x3FF. The "csr_addr" field represents the symbolic names that follow address the corresponding CSRs. The "optional_token" field contains an "indirect_ref" parameter that indicates overriding qualifiers or additional qualifiers, fully described below, are associated with this reference.

The 10-bit immediate data for FAST_WRT includes the following fields. The "incr_enq_num1" and "incr_enq_num2" fields write a 1 to increment an enqueue sequence number by one. The enqueue sequence number is always incremented by one.

5 The "inter_thd_sig" field represents a thread number (0 - 23) of the thread to be signaled.

The 10-bit immediate data supplied with the instruction is shifted in two segments to the appropriate fields. Bits 6 through 0 are shifted left by an amount equal to the thread number writing the data. Bits 9 through 7 are always shifted into the BP2 through BP0 positions regardless of the micro engine writing the data.

10 The "self_destruct" field specifies a bit position (0 to 31) that will be set. A "thread_done" field represents a 2-bit message that is shifted into a position relative to the thread that is writing the message. The message is determined through software.

The "thread_done_incr1" and "thread_done_incr2" fields are the same as for the thread_done register except that either "enqueue_seq1" or "enqueue_seq2" is
15 incremented.

The "xmit_validate" field represents the transmit FIFO element number (0 to 15) that is marked to indicate that the data is valid in this element.

The "indirect_ref" field has the bits definitions and functions set out below:

20	Bits	Field	Description
	31	OV	If set, the UENG ADDR field overrides the default micro engine address implied by the FAST_WR command, which is the micro engine that issued the CSR reference.
25	30:28	UENG ADDR	Specifies the micro engine associated with the CSR reference. If bit [31] = 0, this field is ignored. Valid UENG ADDR values are 0 through 5.
	27:16	RES	Reserved. Returns 0 when read.
30	15	OV	If set, the FAST WRITE DATA field overrides the default micro engine data specified by fast_wr_data.

	14:5	FASTWRITEDATA	Immediate data to be written to the CSR. Valid FAST WRITE DATA values are 0 through 0x3FF. If bit [15] = 0, this field is ignored.
	4:3	RES	Reserved. Returns 0 when read.
5	2	OV	If set, the CTX field overrides the default context implied by the FAST_WR command.
	1:0	CTX	Specifies the context associated with the CSR reference. If bit [2] = 0, this field is ignored. Valid CTX values are 0 through 3.
10			

It is to be understood that while the invention has been described in conjunction with the detailed description thereof, the foregoing description is intended to illustrate and not limit the scope of the invention, which is defined by the scope of the appended claims.

15 Other aspects, advantages, and modifications are within the scope of the following claims.

WHAT IS CLAIMED IS:

1. A method of operating a processor comprising:
receiving data in a processing thread having a processing thread number; and
loading the data into a register having a register address corresponding to the
processing thread number.
- 5 2. The method of claim 1 wherein the register is a control and status register (CSR).
3. The method of claim 2 wherein the status and control register is contained in a 64-
bit wide first-in first-out (FIFO) bus.
4. The method of claim 3 wherein the FIFO bus interfaces with Media Access
Controller (MAC) devices.
- 10 5. The method of claim 1 wherein the data represents values hexadecimal mask 0 to
hexadecimal mask 0x3FF.
6. The method of claim 1 wherein the processing thread represents processing in a
micro engine of a multi-threaded processor.
7. The method of claim 1 shifting comprises:
15 shifting a first portion of the data by an amount equal to the processing thread
number; and
shifting second portion of the data into a breakpoint (BP) register 2 to a BP
register 0.
8. The method of claim 1 wherein receiving the data further comprises receiving a
20 token.
9. The method of claim 8 wherein the token represents overriding qualifiers.
10. The method of claim 8 wherein the token is a 32-bit word.

11. The method of claim 10 wherein a token format comprises:
an OV field in bit 31;
an micro engine (UENG) ADDR field in bits 30:28;
a reserved field in bits 27:16;
5 an OV field in bit 15;
a fast write data field in bits 14:5;
a reserved field in bits 4:3;
an OV field in bit 2; and
a CTX field in bits 1:0.
- 10 12. The method of claim 11 wherein a micro engine address overrides a default micro engine address if bit 31 is set.
13. The method of claim 11 wherein bits 30:28 specify a micro engine associated with a control and status register (CSR).
14. The method of claim 11 wherein bits 27:16 return 0 when read.
- 15 15. The method of claim 11 wherein a micro engine address overrides a default micro engine address if bit 15 is set.
16. The method of claim 11 wherein bits 14:5 represent valid data to be written to a control and status register (CSR).
17. The method of claim 11 wherein bits 4:3 return 0 when read.
- 20 18. The method of claim 11 wherein a context (CTX) field overrides a default context if bit 2 is set.
19. The method of claim 11 wherein bits 1:0 specify a context associated with a control and status register (CSR) reference.

20. A computer program product, disposed on a computer readable medium, the program comprising instructions for causing a computer to:

receive data in a processing thread having a processing thread number; and
load the data into a register corresponding to the processing thread number.

5 21. The computer program product of claim 20 wherein the register field is a control and status register (CSR).

22. The computer program product of claim 21 wherein the status and control register is contained in a 64-bit wide first-in first-out (FIFO) bus.

10 23. The computer program product of claim 22 wherein the FIFO bus interfaces with Media Access Controller (MAC) devices.

24. The computer program product of claim 20 wherein the data represents values hexadecimal mask 0 to hexadecimal mask 0x3FF.

25. The computer program product of claim 20 wherein the processing thread represents processing in a micro engine of a multi-threaded processor.

15 26. The computer program product of claim 20 further comprising instructions for causing the computer to:
shift a first portion of the data left by an amount equal to the processing thread number; and
shift a second portion of the data into a breakpoint (BP) register 2 to a BP register
20 0.

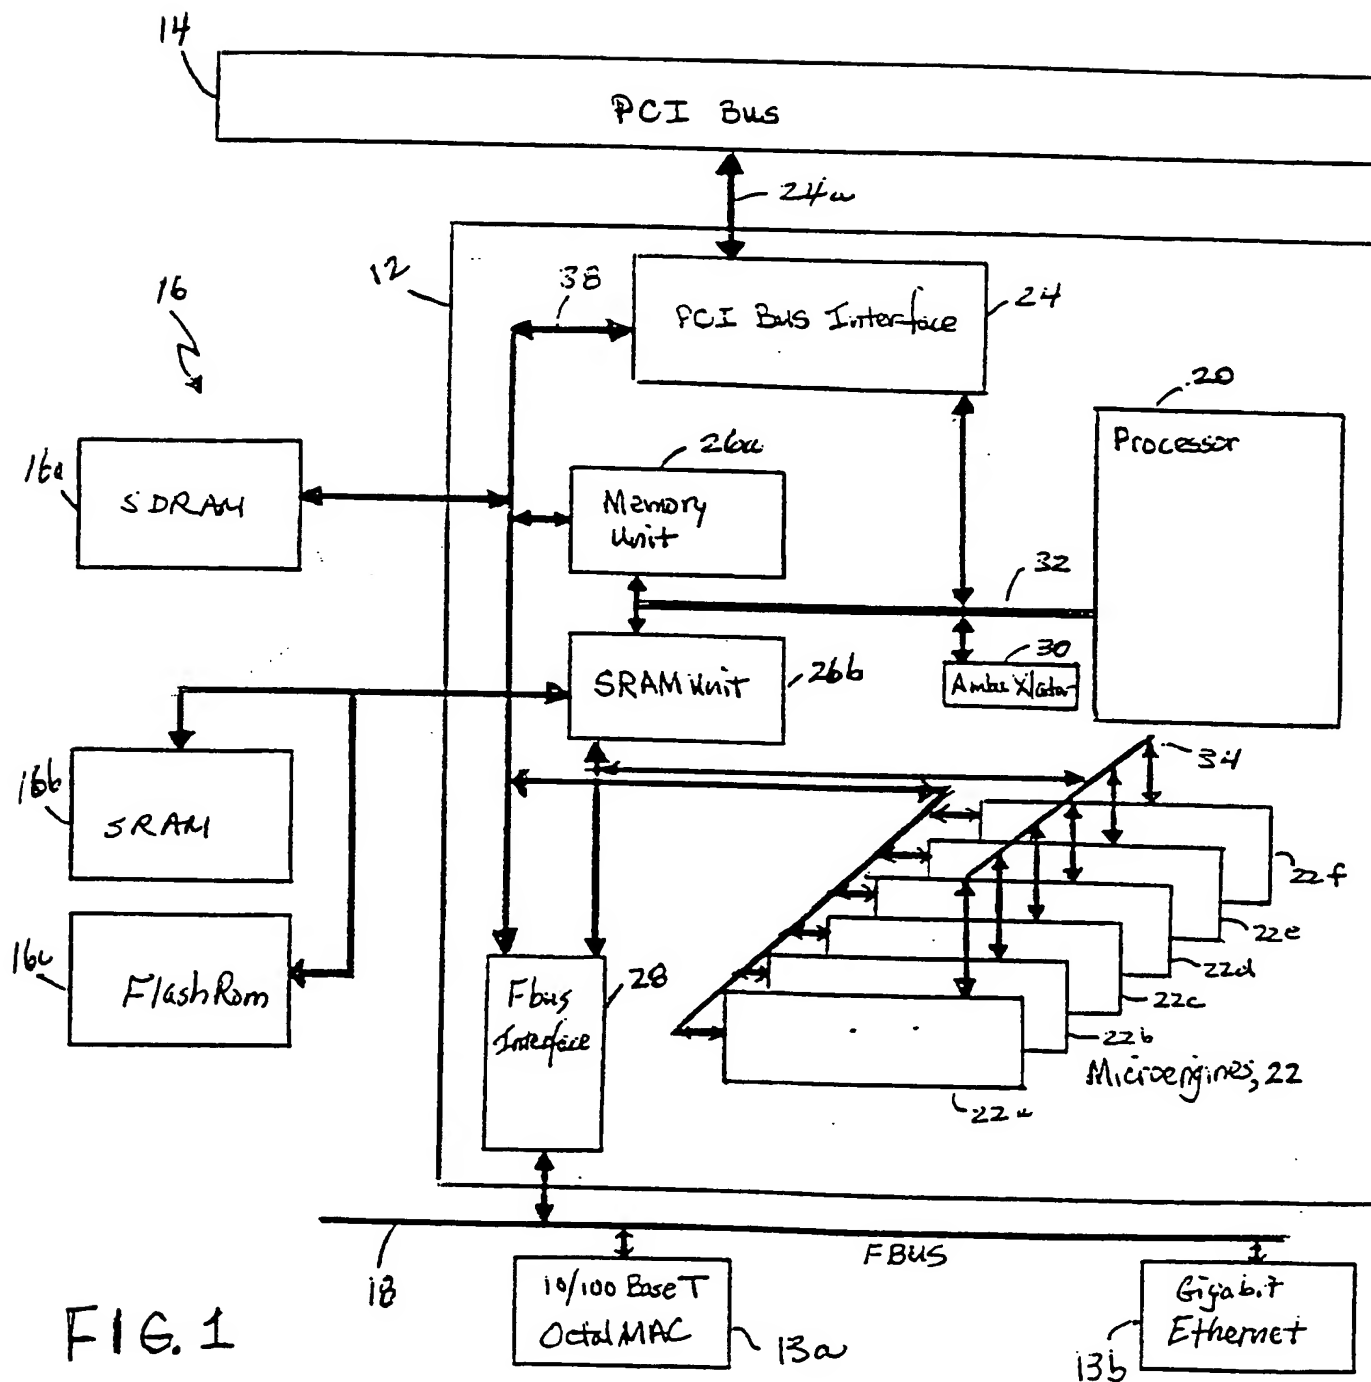
27. The computer program product of claim 20 further comprising an instruction for causing the computer to receive a token.

28. The computer program product of claim 27 wherein the token represents overriding qualifiers.

25 29. The computer program product of claim 27 wherein the token is a 32-bit word.

30. The computer program product of claim 29 wherein a token format comprises:
an OV field in bit 31;
an micro engine (UENG) ADDR field in bits 30:28;
a reserved field in bits 27:16;
an OV field in bit 15;
a fast write data field in bits 14:5;
a reserved field in bits 4:3;
an OV field in bit 2; and
a CTX field in bits 1:0.
31. The computer program product of claim 30 wherein a micro engine address overrides a default micro engine address if bit 31 is set.
32. The computer program product of claim 30 wherein bits 30:28 specify a micro engine associated with a control and status register (CSR).
33. The computer program product of claim 30 wherein bits 27:16 return 0 when read.
34. The computer program product of claim 30 wherein a micro engine address overrides a default micro engine address if bit 15 is set.
35. The computer program product of claim 30 wherein bits 14:5 represent valid data to be written to a control and status register (CSR).
36. The computer program product of claim 30 wherein bits 4:3 return 0 when read.
37. The computer program product of claim 30 a context (CTX) field overrides a default context if bit 2 is set.
38. The computer program product of claim 30 wherein bits 1:0 specify a context associated with a control and status register (CSR) reference.
- ABSTRACT
A method of operating a processor including receiving data in a processing thread having a processing thread number and shifting the data into a register corresponding to the processing thread number.

10



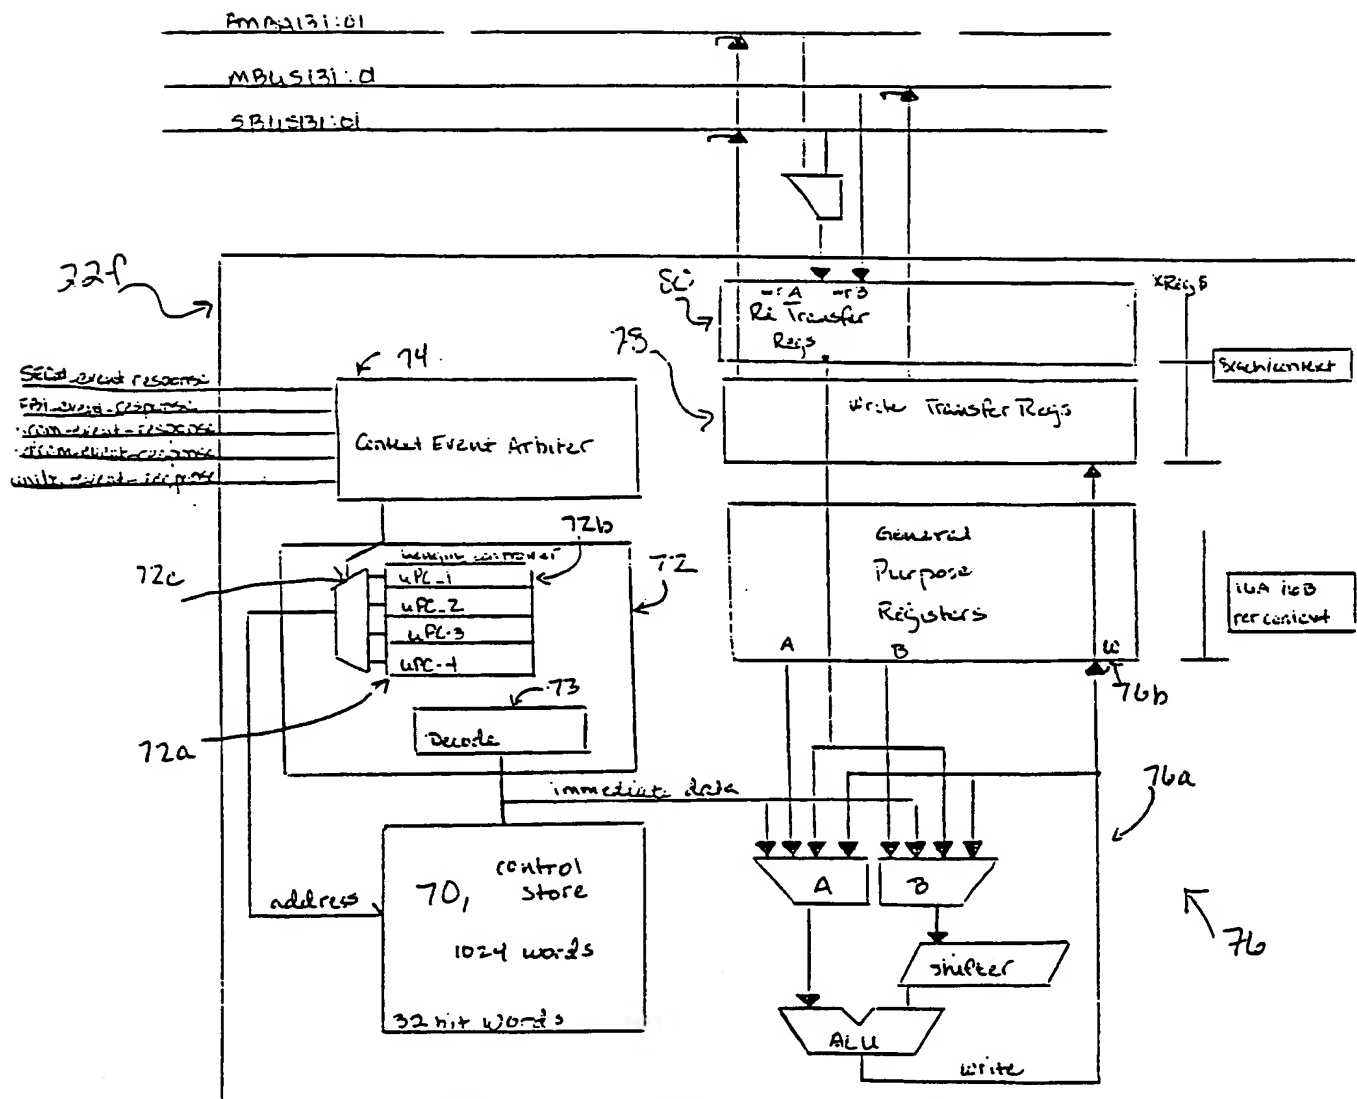


Fig 3

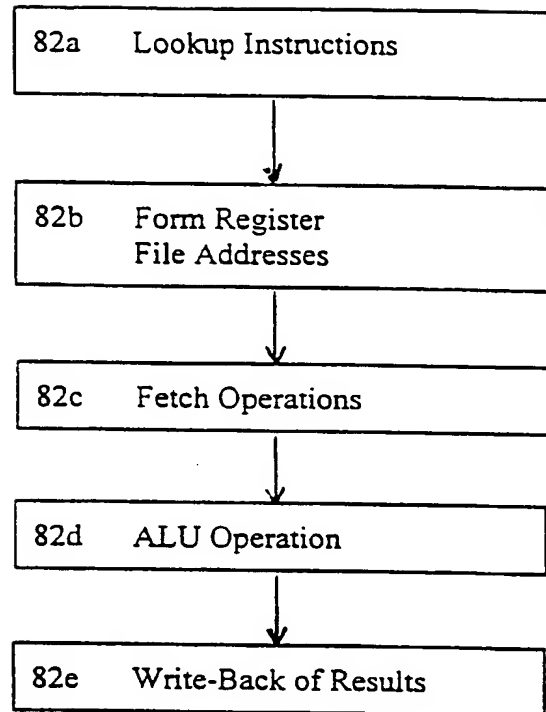
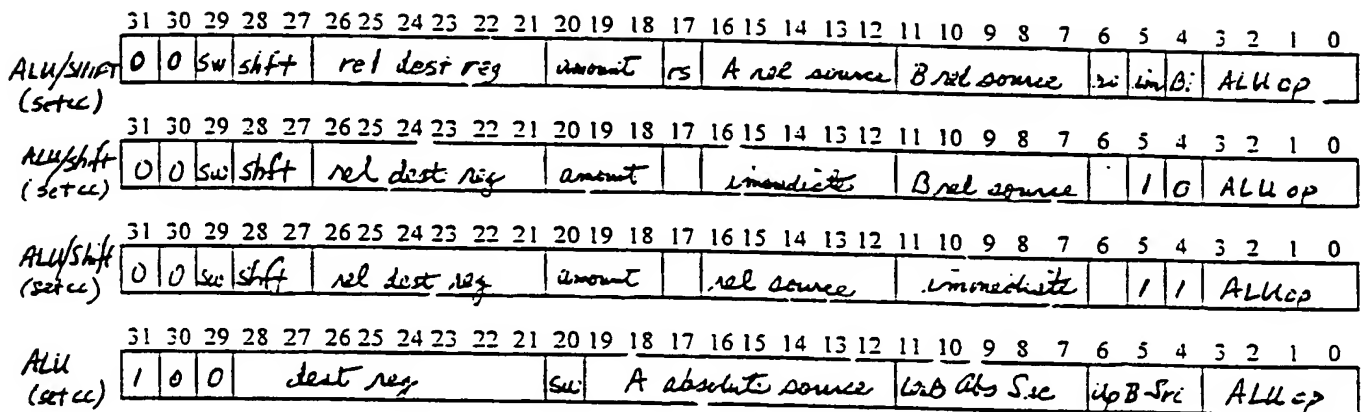


FIG. 4



Shift Decode:

(rs, r0) decode ([31:0] shifts into [63:32] and take [63:32]):

00 = left rotate

01 = right shift (32-ShfAmt = Right Shift Amt)

10 = left shift

11 = double shift (upper A-op shifts into lower B-op)

==> "left rotate" of zero gives zero shift (otherwise zero amount signifies indirect shift)

ALU-OP decode:

0000 = B

0001 = ~B

0010 = A&B (and)

0011 = A&~B (and~)

0100 = ~A&B (~and)

0101 = XOR

0110 = OR

0111 = mul-stuff

1000 = A-B

1001 = B-A

1010 =

1011 =

1100 = A+B(8)

1101 = A+B(16)

1110 = A+B

1111 = A+B+Cin

FIG. 5

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US00/24000**A. CLASSIFICATION OF SUBJECT MATTER**

IPC(7) :G06F 9/312, 9/315, 9/38

US CL :712/43, 224, 225, 228, 229

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 712/43, 224, 225, 228, 229

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

USPAT, EPO, JPO, ACM, IEEE, Derwent, IBM Technical Disclosure Bulletins

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X,P	US 6,058,465 A (NGUYEN) 02 May 2000, col, 39-40 and 91-92.	1-2, 5-6, 8-10, 20-21, 24-25, and 27-29
A	US 6,002,881 A (YORK et al.) 14 December 1999	1-38
A	Waldspurger et al. Register Relocation: Flexible Contexts for Multithreading Proceedings of the 20th Annual International Symposium on Computer Architecture 1993 pages 120-130.	1-38

☐ Further documents are listed in the continuation of Box C.☐ See patent family annex.

* Special categories of cited documents:	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"A" document defining the general state of the art which is not considered to be of particular relevance	"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"E" earlier document published on or after the international filing date	"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"&" document member of the same patent family
"O" document referring to an oral disclosure, use, exhibition or other means	
"P" document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search

14 DECEMBER 2000

Date of mailing of the international search report

08 JAN 2001

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer

WILLIAM M. KRAI

Telephone No. (703) 305-9699

CORRECTED VERSION

(19) World Intellectual Property Organization
International Bureau



20069,352

(43) International Publication Date
8 March 2001 (08.03.2001)

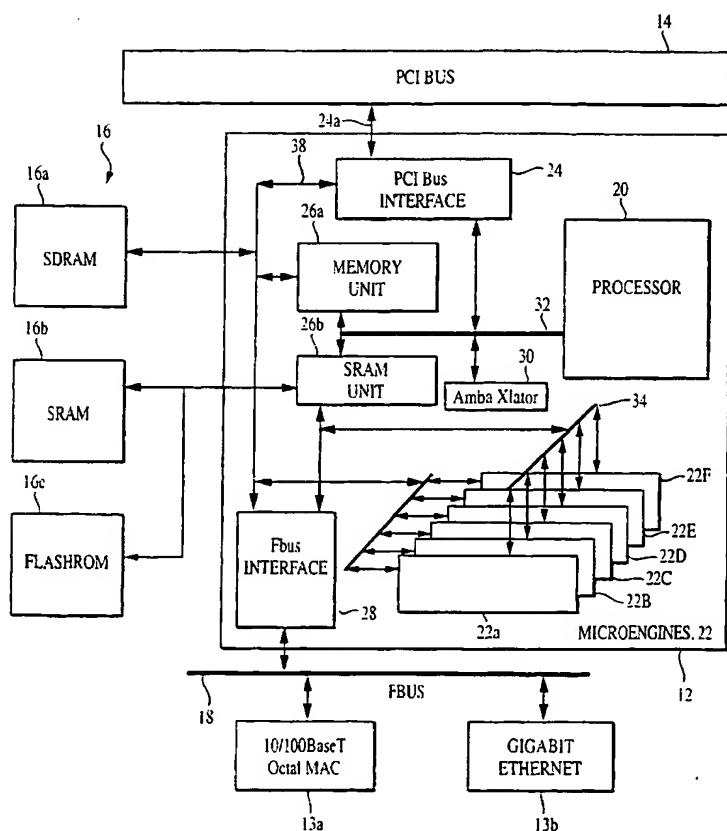
PCT

(10) International Publication Number
WO 01/016714 A1

- (51) International Patent Classification⁷: **G06F 9/312**, 9/315, 9/38
- (21) International Application Number: **PCT/US00/24000**
- (22) International Filing Date: **31 August 2000 (31.08.2000)**
- (25) Filing Language: **English**
- (26) Publication Language: **English**
- (30) Priority Data:
60/151,961 1 September 1999 (01.09.1999) US
- (63) Related by continuation (CON) or continuation-in-part (CIP) to earlier application:
US 60/151,961 (CIP)
Filed on 1 September 1999 (01.09.1999)
- (71) Applicant (for all designated States except US): **INTEL CORPORATION** [US/US]; 2200 Mission College Boulevard, Santa Clara, CA 95052 (US).
- (72) Inventors; and
(75) Inventors/Applicants (for US only): **WOLRICH, Gilbert** [US/US]; 4 Cider Mill Road, Framingham, MA 01701 (US). **ADILETTA, Matthew, J.** [US/US]; 20 Monticello Drive, Worcester, MA 01603 (US). **WHEELER, William** [US/US]; 745 School Street, Webster, MA 01570 (US). **BERNSTEIN, Debra** [US/US]; 38 Helen Street, Waltham, MA 02452 (US). **HOOPER, Donald** [US/US]; 19 Main Circle, Shrewsbury, MA 01545 (US).
- (74) Agents: **MALONEY, Denis, G.**; Fish & Richardson P.C., 225 Franklin Street, Boston, MA 02110-2804 et al. (US).

[Continued on next page]

(54) Title: FAST WRITE INSTRUCTION FOR MICRO ENGINE USED IN MULTITHREADED PARALLEL PROCESSOR ARCHITECTURE



(57) Abstract: A method of operating a processor (12) including receiving data in a processing thread having a processing thread number and shifting the data into a register corresponding to the processing thread number.

WO 01/016714 A1



(81) **Designated States (national):** AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW.

(84) **Designated States (regional):** ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

Published:

— with international search report

(48) Date of publication of this corrected version:

12 September 2002

(15) Information about Correction:

see PCT Gazette No. 37/2002 of 12 September 2002, Section II

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

FAST WRITE INSTRUCTION FOR MICRO ENGINE USED IN MULTITHREADED PARALLEL PROCESSOR ARCHITECTURE

5

BACKGROUND

This invention relates to a memory instruction for computer processors.

Parallel processing is an efficient form of information processing of concurrent events in a computing process. Parallel processing demands concurrent execution of many programs in a computer, in contrast to sequential processing. In the context of a parallel processor, parallelism involves doing more than one thing at the same time. Unlike a serial paradigm where all tasks are performed sequentially at a single station or a pipelined machine where tasks are performed at specialized stations, with parallel processing, a number of stations are provided with each capable of performing all tasks. That is, in general all or a number of the stations work simultaneously and independently on the same or common elements of a problem. Certain problems are suitable for solution by applying parallel processing.

DESCRIPTION OF DRAWINGS

The foregoing features and other aspects of the invention will be described further in detail by the accompanying drawings, in which:

20 FIG. 1 is a block diagram of a communication system employing a hardware-based multithreaded processor.

FIG. 2 is a detailed block diagram of the hardware-based multithreaded processor of FIG. 1.

25 FIG. 3 is a block diagram of a micro engine functional unit employed in the hardware-based multithreaded processor of FIGS. 1 and 2.

FIG. 4 is a block diagram of a pipeline in the micro engine of FIG. 3.

FIG. 5 is a block diagram illustrating a format for arithmetic logic unit instruction results.

Like reference symbols in the various drawings indicate like elements.

DETAILED DESCRIPTION

Referring to FIG. 1, a communication system 10 includes a parallel, hardware-based multithreaded processor 12. The hardware-based multithreaded processor 12 is coupled to a bus such as a PCI bus 14, a memory system 16 and a second bus 18. The system 10 is especially useful for tasks that can be broken into parallel subtasks or functions. Specifically, hardware-based multithreaded processor 12 is useful for tasks that are bandwidth oriented rather than latency oriented. The hardware-based multithreaded processor 12 has multiple micro engines 22 each with multiple hardware controlled threads that can be simultaneously active and independently work on a task.

The hardware-based multithreaded processor 12 also includes a central controller 20 that assists in loading micro code control for other resources of the hardware-based multithreaded processor 12 and performs other general purpose computer type functions such as handling protocols, exceptions, extra support for packet processing where the micro engines 22 pass the packets off for more detailed processing such as in boundary conditions. In one embodiment, the processor 20 is a Strong Arm® (Arm is a trademark of ARM Limited, United Kingdom) based architecture. The general-purpose microprocessor 20 has an operating system. Through the operating system the processor 20 can call functions to operate on micro engines 22a-22f. The processor 20 can use any supported operating system, preferably a real time operating system. For the core processor 20 implemented as Strong Arm architecture, operating systems such as, Microsoft-NT real-time, VXWorks and μ CUS, a freeware operating system available over the Internet, can be used.

Functional micro engines (micro engines) 22a-22f each maintain program counters in hardware and states associated with the program counters. Effectively, a corresponding number of sets of threads can be simultaneously active on each of the micro engines 22a-22f while only one is actually operating at any one time.

In an embodiment, there are six micro engines 22a-22f as shown. Each micro engine 22a-22f has capabilities for processing four hardware threads. The six micro engines 22a-22f operate with shared resources including memory system 16 and bus interfaces 24 and 28. The memory system 16 includes a Synchronous Dynamic Random Access Memory (SDRAM) controller 26a and a Static Random Access Memory (SRAM) controller 26b. SDRAM memory 16a and SDRAM controller 26a are typically used for processing large volumes of data, e.g., processing of network payloads from network

packets. The SRAM controller 26b and SRAM memory 16b are used in a networking implementation for low latency, fast access tasks, e.g., accessing look-up tables, memory for the core processor 20, and so forth.

5 The six micro engines 22a-22f access either the SDRAM 16a or SRAM 16b based on characteristics of the data. Thus, low latency, low bandwidth data is stored in and fetched from SRAM 16b, whereas higher bandwidth data for which latency is not as important, is stored in and fetched from SDRAM 16a. The micro engines 22a-22f can execute memory reference instructions to either the SDRAM controller 26a or SRAM controller 16b.

10 Advantages of hardware multithreading can be explained by SRAM or SDRAM memory accesses. As an example, an SRAM access requested by a Thread_0, from a micro engine will cause the SRAM controller 26b to initiate an access to the SRAM memory 16b. The SRAM controller 26b controls arbitration for the SRAM bus, accesses the SRAM 16b, fetches the data from the SRAM 16b, and returns data to a requesting
15 micro engine 22a-22f. During an SRAM access, if the micro engine, e.g., micro engine 22a, had only a single thread that could operate, that micro engine would be dormant until data was returned from the SRAM 16b. By employing hardware context swapping within each of the micro engines 22a-22f, the hardware context swapping enables other contexts with unique program counters to execute in that same micro engine. Thus, another
20 thread, e.g., Thread_1 can function while the first thread, i.e., Thread_0, is awaiting the read data to return. During execution, Thread_1 may access the SDRAM memory 16a. While Thread_1 operates on the SDRAM unit 16a, and Thread_0 is operating on the SRAM unit 16b, a new thread, e.g., Thread_2 can now operate in the micro engine 22a. Thread_2 can operate for a certain amount of time until it needs to access memory or
25 perform some other long latency operation, such as making an access to a bus interface. Therefore, simultaneously, the processor 12 can have a bus operation, SRAM operation and SDRAM operation all being completed or operated upon by one micro engine 22a and have one more thread available to process more work in the data path.

The hardware context swapping also synchronizes completion of tasks. For
30 example, two threads could hit the same shared resource e.g., SRAM 16b. Each one of these separate functional units, e.g., the FBUS interface 28, the SRAM controller 26a, and the SDRAM controller 26b, when they complete a requested task from one of the micro engine thread contexts reports back a flag signaling completion of an operation. When

the micro engine receives the flag, the micro engine can determine which thread to turn on.

An application for the hardware-based multithreaded processor 12 is as a network processor. As a network processor, the hardware-based multithreaded processor 12
5 interfaces to network devices such as a media access controller device e.g., a 10/100BaseT Octal MAC 13a or a Gigabit Ethernet device 13b. In general, as a network processor, the hardware-based multithreaded processor 12 can interface to any type of communication device or interface that receives/sends large amounts of data. Communication system 10 functioning in a networking application could receive network
10 packets from the devices 13a, 13b and process those packets in a parallel manner. With the hardware-based multithreaded processor 12, each network packet can be independently processed.

Another example for use of processor 12 is a print engine for a postscript processor or as a processor for a storage subsystem, e.g., Redundant Array of Independent
15 Disk (RAID) storage, a category of disk drives that employs two or more drives in combination for fault tolerance and performance. A further use is as a matching engine. In the securities industry for example, the advent of electronic trading requires the use of electronic matching engines to match orders between buyers and sellers. These and other parallel types of tasks can be accomplished utilizing the system 10.

20 The processor 12 includes the bus interface 28 that couples the processor to the second bus 18. In an embodiment, bus interface 28 couples the processor 12 to the FBUS (FIFO bus) 18. The FBUS interface 28 is responsible for controlling and interfacing the processor 12 to the FBUS 18. The FBUS 18 is a 64-bit wide FIFO bus, used to interface to Media Access Controller (MAC) devices, e.g., 10/100 Base T Octal MAC 13a.

25 The processor 12 includes a second interface e.g., PCI bus interface 24, that couples other system components that reside on the PCI 14 bus to the processor 12. The PCI bus interface 24 provides a high-speed data path 24a to memory 16, e.g., SDRAM memory 16a. Through PCI bus interface 24 data can be moved quickly from the SDRAM 16a through the PCI bus 14, via direct memory access (DMA) transfers. The hardware
30 based multithreaded processor 12 supports image transfers. The hardware based multithreaded processor 12 can employ DMA channels so if one target of a DMA transfer is busy, another one of the DMA channels can take over the PCI bus 14 to deliver information to another target to maintain high processor 12 efficiency. Additionally, the

PCI bus interface 24 supports target and master operations. Target operations are operations where slave devices on bus 14 access SDRAMs through reads and writes that are serviced as a slave to a target operation. In master operations, the processor core 20 sends data directly to or receives data directly from the PCI interface 24.

5 Each of the functional units 22 is coupled to one or more internal buses. As described below, the internal buses are dual, 32 bit buses (i.e., one bus for read and one for write). The hardware-based multithreaded processor 12 also is constructed such that the sum of the bandwidths of the internal buses in the processor 12 exceed the bandwidth of external buses coupled to the processor 12. The processor 12 includes an internal core
10 processor bus 32, e.g., an ASB Advanced System Bus (ASB), that couples the processor core 20 to the memory controller 26a, 26b and to an ASB translator 30, described below. The ASB bus 32 is a subset of the so-called Advanced Microcontroller Bus Architecture (AMBA) bus that is used with the Strong Arm processor core 20. AMBA is an open
15 standard, on-chip bus specification that details a strategy for the interconnection and management of functional blocks that makes up a System-on-chip (SoC). The processor 12 also includes a private bus 34 that couples the micro engine units 22 to SRAM
20 controller 26b, ASB translator 30 and FBUS interface 28. A memory bus 38 couples the memory controller 26a, 26b to the bus interfaces 24 and 28 and memory system 16 including flashrom 16c that is used for boot operations and so forth.

20 Referring to FIG. 2, each of the micro engines 22a-22f includes an arbiter that examines flags to determine the available threads to be operated upon. Any thread from any of the micro engines 22a-22f can access the SDRAM controller 26a, SDRAM
25 controller 26b or FBUS interface 28. The memory controllers 26a and 26b each include queues to store outstanding memory reference requests. The queues either maintain order of memory references or arrange memory references to optimize memory bandwidth. For
30 example, if a thread_0 has no dependencies or relationship to a thread_1, there is no reason that thread_1 and thread_0 cannot complete their memory references to the SRAM unit 16b out of order. The micro engines 22a-22f issue memory reference requests to the memory controllers 26a and 26b. The micro engines 22a-22f flood the memory
30 subsystems 26a and 26b with enough memory reference operations such that the memory subsystems 26a and 26b become the bottleneck for processor 12 operation.

 If the memory subsystem 16 is flooded with memory requests that are independent in nature, the processor 12 can perform memory reference sorting. Memory reference

sorting improves achievable memory bandwidth. Memory reference sorting, as described below, reduces dead time or a bubble that occurs with accesses to SRAM 16b. With memory references to SRAM 16b, switching current direction on signal lines between reads and writes produces a bubble or a dead time waiting for current to settle on
 5 conductors coupling the SRAM 16b to the SRAM controller 26b.

That is, the drivers that drive current on the bus need to settle out prior to changing states. Thus, repetitive cycles of a read followed by a write can degrade peak bandwidth. Memory reference sorting allows the processor 12 to organize references to memory such that long strings of reads can be followed by long strings of writes. This
 10 can be used to minimize dead time in the pipeline to effectively achieve closer to maximum available bandwidth. Reference sorting helps maintain parallel hardware context threads. On the SDRAM 16a, reference sorting allows hiding of pre-charges from one bank to another bank. Specifically, if the memory system 16b is organized into an odd bank and an even bank, while the processor is operating on the odd bank, the memory
 15 controller can start pre-charging the even bank. Pre-charging is possible if memory references alternate between odd and even banks. By ordering memory references to alternate accesses to opposite banks, the processor 12 improves SDRAM bandwidth. Additionally, other optimizations can be used. For example, merging optimizations where operations that can be merged, are merged prior to memory access, open page
 20 optimizations where by examining addresses an opened page of memory is not reopened, chaining, as will be described below, and refreshing mechanisms, can be employed.

The FBUS interface 28 supports Transmit and Receive flags for each port that a MAC device supports, along with an Interrupt flag indicating when service is warranted. The FBUS interface 28 also includes a controller 28a that performs header processing of
 25 incoming packets from the FBUS 18. The controller 28a extracts the packet headers and performs a micro programmable source/destination/protocol hashed lookup (used for address smoothing) in SRAM 16b. If the hash does not successfully resolve, the packet header is sent to the processor core 20 for additional processing. The FBUS interface 28 supports the following internal data transactions:

30

FBUS unit	(Shared bus SRAM)	to/from micro engine.
FBUS unit	(via private bus)	writes from SDRAM Unit.
FBUS unit	(via Mbus)	Reads to SDRAM.

The FBUS 18 is a standard industry bus and includes a data bus, e.g., 64 bits wide and sideband control for address and read/write control. The FBUS interface 28 provides the ability to input large amounts of data using a series of input and output FIFOs 29a-29b. From the FIFOs 29a-29b, the micro engines 22a-22f fetch data from or command the SDRAM controller 26a to move data from a receive FIFO in which data has come from a device on bus 18, into the FBUS interface 28. The data can be sent through memory controller 26a to SDRAM memory 16a, via a direct memory access. Similarly, the micro engines can move data from the SDRAM 26a to interface 28, out to FBUS 18, via the FBUS interface 28.

Data functions are distributed amongst the micro engines 22. Connectivity to the SRAM 26a, SDRAM 26b and FBUS 28 is via command requests. A command request can be a memory request or a FBUS request. For example, a command request can move data from a register located in a micro engine 22a to a shared resource, e.g., an SDRAM location, SRAM location, flash memory or some MAC address. The commands are sent out to each of the functional units and the shared resources. However, the shared resources do not need to maintain local buffering of the data. Rather, the shared resources access distributed data located inside of the micro engines 22a-22f. This enables micro engines 22a-22f, to have local access to data rather than arbitrating for access on a bus and risk contention for the bus. With this feature, there is a zero cycle stall for waiting for data internal to the micro engines 22a-22f.

The data buses, e.g., ASB bus 30, SRAM bus 34 and SDRAM bus 38 coupling these shared resources, e.g., memory controllers 26a and 26b, are of sufficient bandwidth such that there are no internal bottlenecks. In order to avoid bottlenecks, the processor 12 has an bandwidth requirement where each of the functional units is provided with at least twice the maximum bandwidth of the internal buses. As an example, the SDRAM 16a can run a 64 bit wide bus at 83 MHz. The SRAM data bus could have separate read and write buses, e.g., could be a read bus of 32 bits wide running at 166 MHz and a write bus of 32 bits wide at 166 MHz. That is, in essence, 64 bits running at 166 MHz which is effectively twice the bandwidth of the SDRAM.

The core processor 20 also can access the shared resources. The core processor 20 has a direct communication to the SDRAM controller 26a to the bus interface 24 and to SRAM controller 26b via bus 32. However, to access the micro engines 22a-22f and

transfer registers located at any of the micro engines 22a-22f, the core processor 20 access the micro engines 22a-22f via the ASB Translator 30 over bus 34. The ASB translator 30 can physically reside in the FBUS interface 28, but logically is distinct. The ASB Translator 30 performs an address translation between FBUS micro engine transfer
 5 register locations and core processor addresses (i.e., ASB bus) so that the core processor 20 can access registers belonging to the micro engines 22a-22f.

Although micro engines 22a-22f can use the register set to exchange data as described below, a scratchpad memory 27 is also provided to permit micro engines 22a-22f to write data out to the memory for other micro engines to read. The scratchpad 27 is
 10 coupled to bus 34.

The processor core 20 includes a RISC core 50 implemented in a five stage pipeline performing a single cycle shift of one operand or two operands in a single cycle, provides multiplication support and 32 bit barrel shift support. This RISC core 50 is a standard Strong Arm® architecture but it is implemented with a five-stage pipeline for
 15 performance reasons. The processor core 20 also includes a 16-kilobyte instruction cache 52, an 8-kilobyte data cache 54 and a prefetch stream buffer 56. The core processor 20 performs arithmetic operations in parallel with memory writes and instruction fetches. The core processor 20 interfaces with other functional units via the ARM defined ASB bus. The ASB bus is a 32-bit bi-directional bus 32.

Referring to FIG. 3, an exemplary one of the micro engines 22a-22f, e.g., micro engine 22f, is shown. The micro engine 22f includes a control store 70, which, in one implementation, includes a RAM of here 1,024 words of 32 bit. The RAM stores a micro program (not shown). The micro program is loadable by the core processor 20. The micro engine 22f also includes controller logic 72. The controller logic 72 includes an
 25 instruction decoder 73 and program counter (PC) units 72a-72d. The four micro program counters 72a-72d are maintained in hardware. The micro engine 22f also includes context event switching logic 74. Context event logic 74 receives messages (e.g., SEQ_#_EVENT_RESPONSE; FBI_EVENT_RESPONSE; SRAM
 _EVENT_RESPONSE; SDRAM_EVENT_RESPONSE; and ASB
 30 _EVENT_RESPONSE) from each one of the shared resources, e.g., SRAM 26a, SDRAM 26b, or processor core 20, control and status registers, and so forth. These messages provide information on whether a requested function has completed. Based on whether or not a function requested by a thread has completed and signaled completion, the thread

needs to wait for that completion signal, and if the thread is enabled to operate, then the thread is placed on an available thread list (not shown). The micro engine 22f can have a maximum of four threads available.

In addition to event signals that are local to an executing thread, the micro engines
5 22a-22f employ signaling states that are global. With signaling states, an executing thread can broadcast a signal state to all micro engines 22a-22f, e.g., Receive Request Available (RRA) signal, any and all threads in the micro engines 22a-22f can branch on these signaling states. These signaling states can be used to determine availability of a resource or whether a resource is due for servicing.

10 The context event logic 74 has arbitration for the four threads. In an embodiment, the arbitration is a round robin mechanism. Other techniques could be used including priority queuing or weighted fair queuing. The micro engine 22f also includes an execution box (EBOX) data path 76 that includes an arithmetic logic unit (ALU) 76a and general-purpose register set 76b. The ALU 76a performs arithmetic and logical functions
15 as well as shift functions. The register set 76b has a relatively large number of general-purpose registers. In an embodiment, there are 64 general-purpose registers in a first bank, Bank A and 64 in a second bank, Bank B. The general-purpose registers are windowed so that they are relatively and absolutely addressable.

The micro engine 22f also includes a write transfer register stack 78 and a read
20 transfer stack 80. These registers 78 and 80 are also windowed so that they are relatively and absolutely addressable. Write transfer register stack 78 is where write data to a resource is located. Similarly, read register stack 80 is for return data from a shared resource. Subsequent to or concurrent with data arrival, an event signal from the respective shared resource e.g., the SRAM controller 26a, SDRAM controller 26b or core
25 processor 20 will be provided to context event arbiter 74, which will then alert the thread that the data is available or has been sent. Both transfer register banks 78 and 80 are connected to the execution box (EBOX) 76 through a data path. In an embodiment, the read transfer register has 64 registers and the write transfer register has 64 registers.

Referring to FIG. 4, the micro engine data path maintains a 5-stage micro-pipeline
30 82. This pipeline includes lookup of microinstruction words 82a, formation of the register file addresses 82b, read of operands from register file 82c, ALU shift or compare operations 82d, and write-back of results to registers 82e. By providing a write-back data bypass into the ALU/shifter units, and by assuming the registers are implemented as a

register file (rather than a RAM), the micro engine 22f can perform a simultaneous register file read and write, which completely hides the write operation.

The SDRAM interface 26a provides a signal back to the requesting micro engine on reads that indicates whether a parity error occurred on the read request. The micro engine micro code is responsible for checking the SDRAM 16a read Parity flag when the micro engine uses any return data. Upon checking the flag, if it was set, the act of branching on it clears it. The Parity flag is only sent when the SDRAM 16a is enabled for checking, and the SDRAM 16a is parity protected. The micro engines 22 and the PCI Unit 14 are the only requestors notified of parity errors. Therefore, if the processor core 20 or FIFO 18 requires parity protection, a micro engine assists in the request. The micro engines 22a-22f support conditional branches.

Referring to FIG. 5, a format for arithmetic logic unit (ALU) instruction is shown. The micro engines 22 support an instruction set. The instruction set includes instructions that perform an operation on one or two operands and deposit the result into the destination register, and update ALU condition codes according to the result of the operation. Condition codes are lost during context swaps.

The computer instruction architecture set also includes a fast write (FAST_WR) instruction. The FAST_WR instruction for immediate data writes specified immediate data to a specified control and status register (CSR) by having write data specified directly in the instruction rather than in a transfer register. FAST_WR writes the specified immediate data to the specified FBI CSR. This improves performance by eliminating the need for the FBI unit to pull the data from a transfer register. The FBI unit automatically shifts the immediate data into the appropriate register field corresponding to the thread that is writing the FAST_WR data.

A format for the fast write instruction is: fast_wr [immed_data, csr_addr], optional_token, where a description of each of the fields follows.

The "immed_data" field represents 10 bits of the immediate data to be written to the control and status register (CSR); valid immed_data values are 0 through 0x3FF. The "csr_addr" field represents the symbolic names that follow address the corresponding CSRs. The "optional_token" field contains an "indirect_ref" parameter that indicates overriding qualifiers or additional qualifiers, fully described below, are associated with this reference.

The 10-bit immediate data for FAST_WRT includes the following fields. The "incr_enq_num1" and "incr_enq_num2" fields write a 1 to increment an enqueue sequence number by one. The enqueue sequence number is always incremented by one.

The "inter_thd_sig" field represents a thread number (0 - 23) of the thread to be signaled.

The 10-bit immediate data supplied with the instruction is shifted in two segments to the appropriate fields. Bits 6 through 0 are shifted left by an amount equal to the thread number writing the data. Bits 9 through 7 are always shifted into the BP2 through BP0 positions regardless of the micro engine writing the data.

The "self_destruct" field specifies a bit position (0 to 31) that will be set. A "thread_done" field represents a 2-bit message that is shifted into a position relative to the thread that is writing the message. The message is determined through software.

The "thread_done_incr1" and "thread_done_incr2" fields are the same as for the thread_done register except that either "enqueue_seq1" or "enqueue_seq2" is incremented.

The "xmit_validate" field represents the transmit FIFO element number (0 to 15) that is marked to indicate that the data is valid in this element.

The "indirect_ref" field has the bits definitions and functions set out below:

20	Bits	Field	Description
	31	OV	If set, the UENG ADDR field overrides the default micro engine address implied by the FAST_WR command, which is the micro engine that issued the CSR reference.
25	30:28	UENG ADDR	Specifies the micro engine associated with the CSR reference. If bit [31] = 0, this field is ignored. Valid UENG ADDR values are 0 through 5.
	27:16	RES	Reserved. Returns 0 when read.
30	15	OV	If set, the FAST WRITE DATA field overrides the default micro engine data specified by fast_wr_data.

	14:5	FASTWRITEDATA	Immediate data to be written to the CSR. Valid FAST WRITE DATA values are 0 through 0x3FF. If bit [15] = 0, this field is ignored.
	4:3	RES	Reserved. Returns 0 when read.
5	2	OV	If set, the CTX field overrides the default context implied by the FAST_WR command.
	1:0	CTX	Specifies the context associated with the CSR reference. If bit [2] = 0, this field is ignored. Valid CTX values are 0 through 3.
10			

It is to be understood that while the invention has been described in conjunction with the detailed description thereof, the foregoing description is intended to illustrate and not limit the scope of the invention, which is defined by the scope of the appended claims.

15 Other aspects, advantages, and modifications are within the scope of the following claims.

WHAT IS CLAIMED IS:

1. A method of operating a processor comprising:
receiving data in a processing thread having a processing thread number; and
loading the data into a register having a register address corresponding to the
processing thread number.
- 5 2. The method of claim 1 wherein the register is a control and status register (CSR).
3. The method of claim 2 wherein the status and control register is contained in a 64-bit wide first-in first-out (FIFO) bus.
4. The method of claim 3 wherein the FIFO bus interfaces with Media Access Controller (MAC) devices.
- 10 5. The method of claim 1 wherein the data represents values hexadecimal mask 0 to hexadecimal mask 0x3FF.
6. The method of claim 1 wherein the processing thread represents processing in a micro engine of a multi-threaded processor.
7. The method of claim 1 shifting comprises:
15 shifting a first portion of the data by an amount equal to the processing thread number; and
shifting second portion of the data into a breakpoint (BP) register 2 to a BP register 0.
8. The method of claim 1 wherein receiving the data further comprises receiving a
20 token.
9. The method of claim 8 wherein the token represents overriding qualifiers.
10. The method of claim 8 wherein the token is a 32-bit word.

11. The method of claim 10 wherein a token format comprises:
an OV field in bit 31;
an micro engine (UENG) ADDR field in bits 30:28;
a reserved field in bits 27:16;
5 an OV field in bit 15;
a fast write data field in bits 14:5;
a reserved field in bits 4:3;
an OV field in bit 2; and
a CTX field in bits 1:0.
- 10 12. The method of claim 11 wherein a micro engine address overrides a default micro engine address if bit 31 is set.
13. The method of claim 11 wherein bits 30:28 specify a micro engine associated with a control and status register (CSR).
14. The method of claim 11 wherein bits 27:16 return 0 when read.
- 15 15. The method of claim 11 wherein a micro engine address overrides a default micro engine address if bit 15 is set.
16. The method of claim 11 wherein bits 14:5 represent valid data to be written to a control and status register (CSR).
17. The method of claim 11 wherein bits 4:3 return 0 when read.
- 20 18. The method of claim 11 wherein a context (CTX) field overrides a default context if bit 2 is set.
19. The method of claim 11 wherein bits 1:0 specify a context associated with a control and status register (CSR) reference.

20. A computer program product, disposed on a computer readable medium, the program comprising instructions for causing a computer to:
- receive data in a processing thread having a processing thread number; and
 - load the data into a register corresponding to the processing thread number.
- 5 21. The computer program product of claim 20 wherein the register field is a control and status register (CSR).
22. The computer program product of claim 21 wherein the status and control register is contained in a 64-bit wide first-in first-out (FIFO) bus.
23. The computer program product of claim 22 wherein the FIFO bus interfaces with
- 10 Media Access Controller (MAC) devices.
24. The computer program product of claim 20 wherein the data represents values hexadecimal mask 0 to hexadecimal mask 0x3FF.
25. The computer program product of claim 20 wherein the processing thread represents processing in a micro engine of a multi-threaded processor.
- 15 26. The computer program product of claim 20 further comprising instructions for causing the computer to:
- shift a first portion of the data left by an amount equal to the processing thread number; and
 - shift a second portion of the data into a breakpoint (BP) register 2 to a BP register
- 20 0.
27. The computer program product of claim 20 further comprising an instruction for causing the computer to receive a token.
28. The computer program product of claim 27 wherein the token represents overriding qualifiers.
- 25 29. The computer program product of claim 27 wherein the token is a 32-bit word.

30. The computer program product of claim 29 wherein a token format comprises:
an OV field in bit 31;
an micro engine (UENG) ADDR field in bits 30:28;
a reserved field in bits 27:16;
5 an OV field in bit 15;
a fast write data field in bits 14:5;
a reserved field in bits 4:3;
an OV field in bit 2; and
a CTX field in bits 1:0.
- 10 31. The computer program product of claim 30 wherein a micro engine address overrides a default micro engine address if bit 31 is set.
32. The computer program product of claim 30 wherein bits 30:28 specify a micro engine associated with a control and status register (CSR).
33. The computer program product of claim 30 wherein bits 27:16 return 0 when read.
- 15 34. The computer program product of claim 30 wherein a micro engine address overrides a default micro engine address if bit 15 is set.
35. The computer program product of claim 30 wherein bits 14:5 represent valid data to be written to a control and status register (CSR).
36. The computer program product of claim 30 wherein bits 4:3 return 0 when read.
- 20 37. The computer program product of claim 30 a context (CTX) field overrides a default context if bit 2 is set.
38. The computer program product of claim 30 wherein bits 1:0 specify a context associated with a control and status register (CSR) reference.
- ABSTRACT
A method of operating a processor including receiving data in a processing thread having
25 a processing thread number and shifting the data into a register corresponding to the processing thread number.

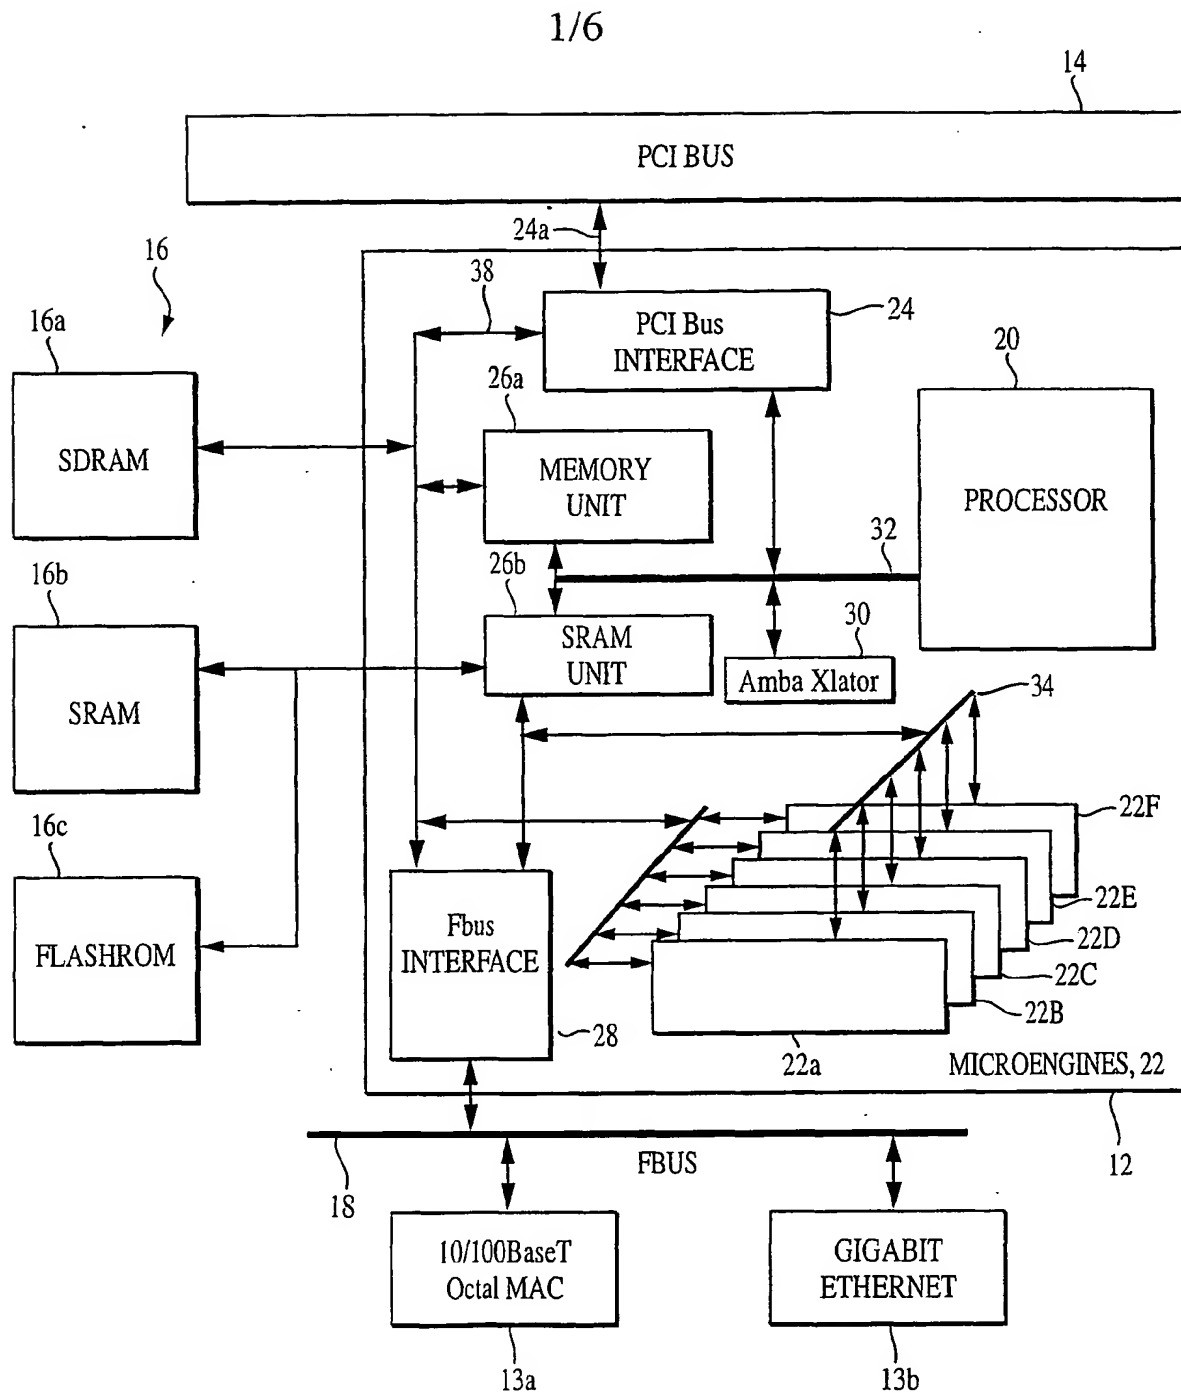


FIG. 1

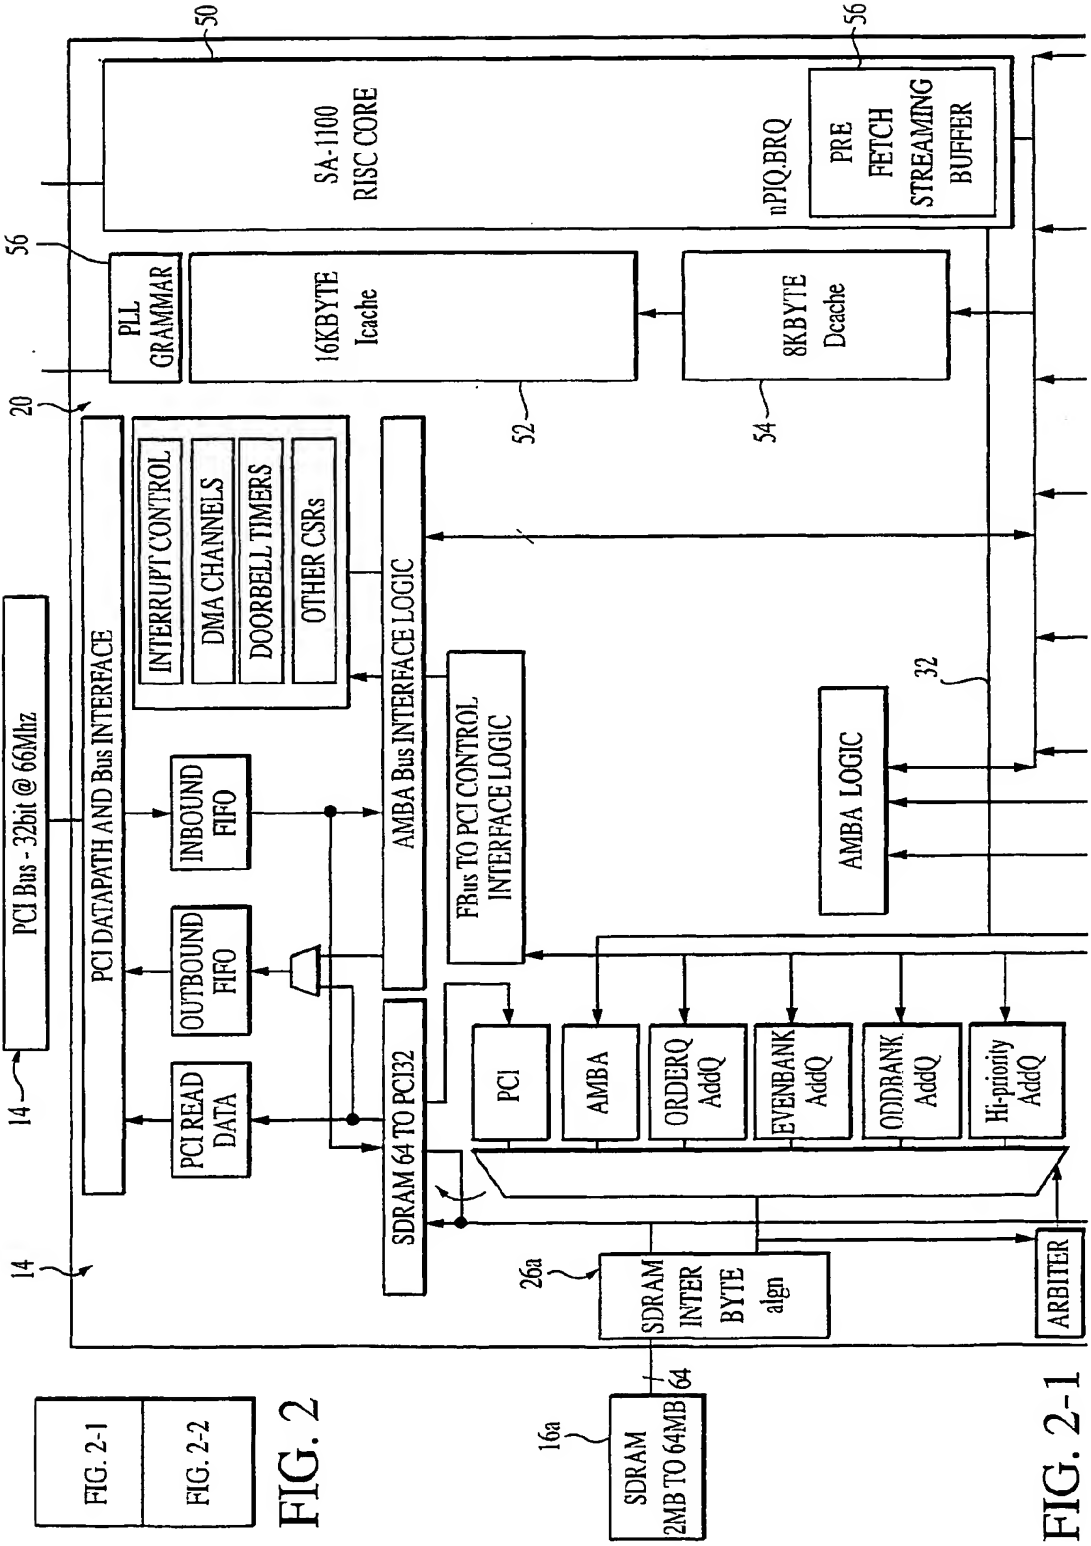


FIG. 2-1

FIG. 2-2

FIG. 2

FIG. 2-1

3/6

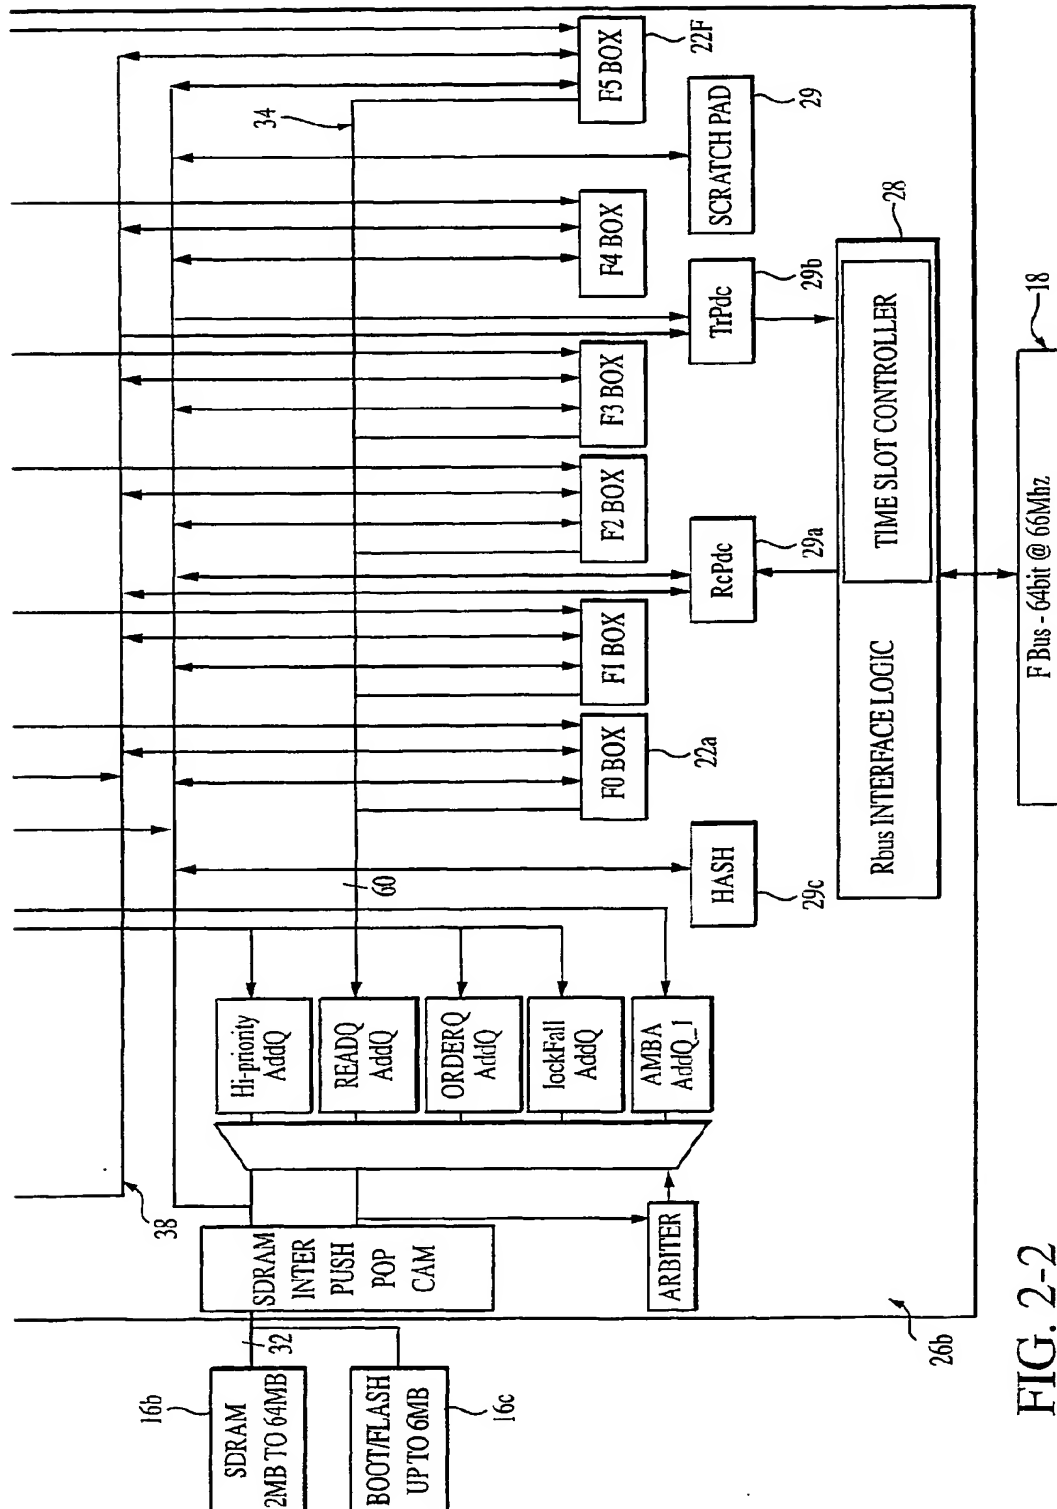


FIG. 2-2

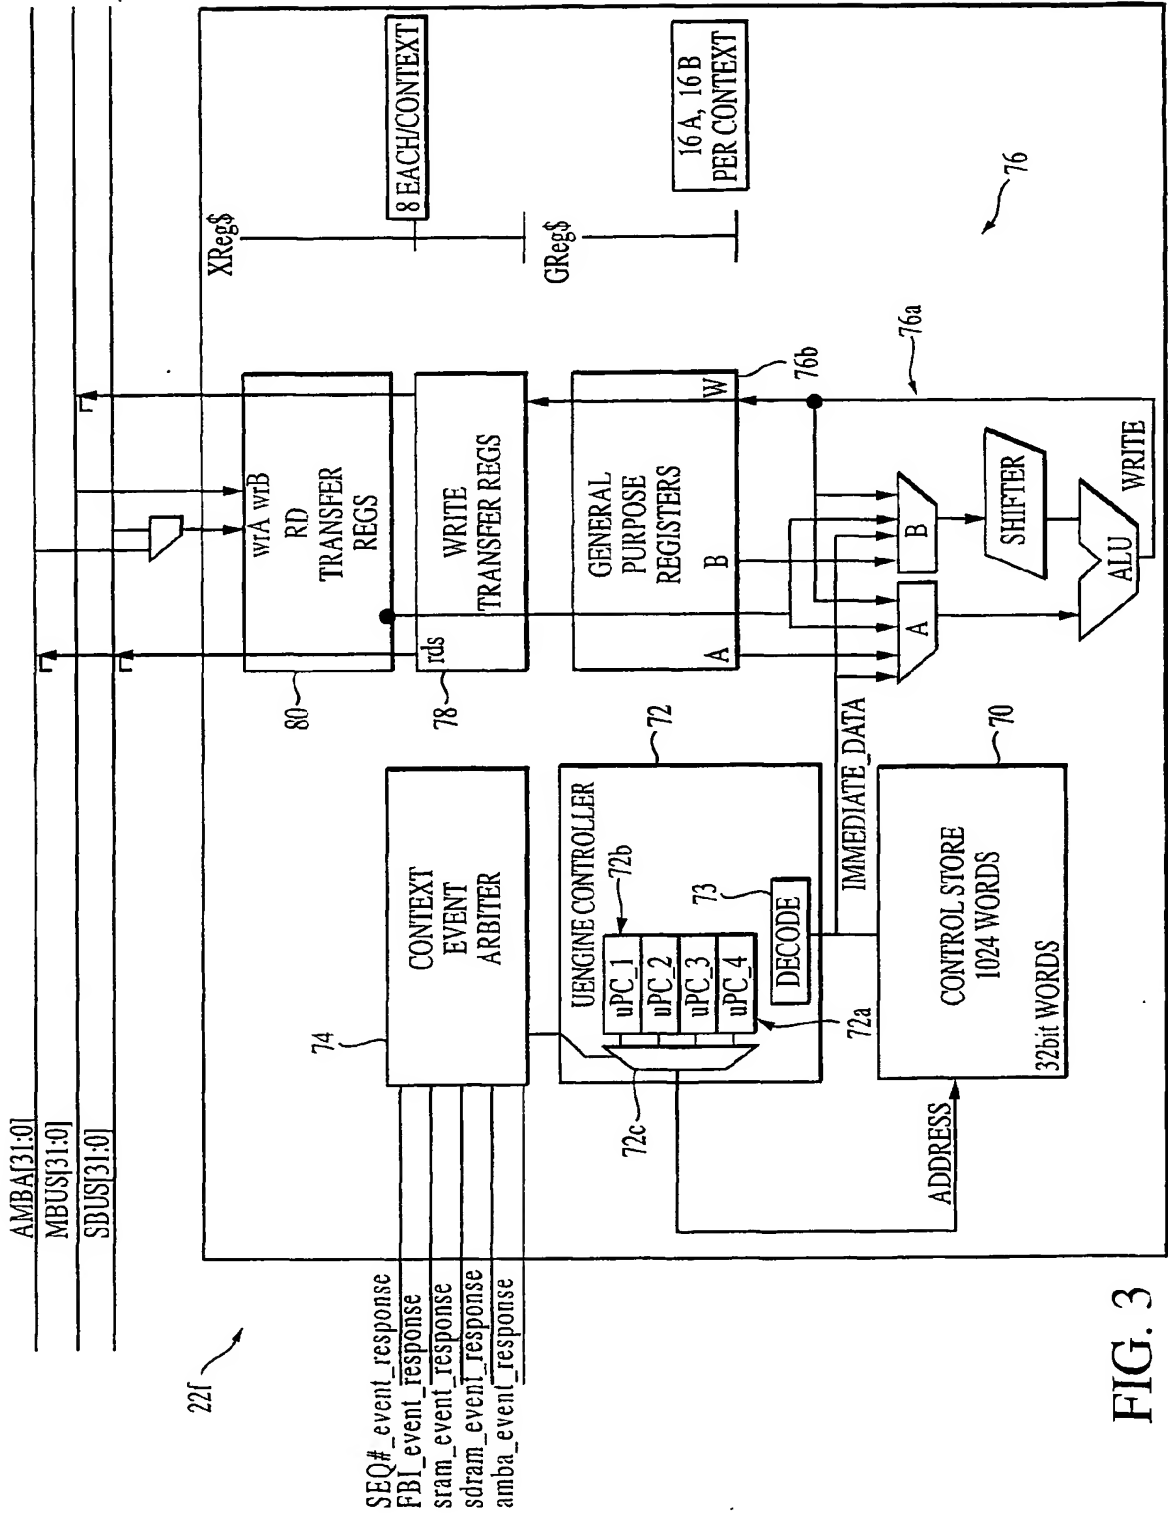


FIG. 3

5/6

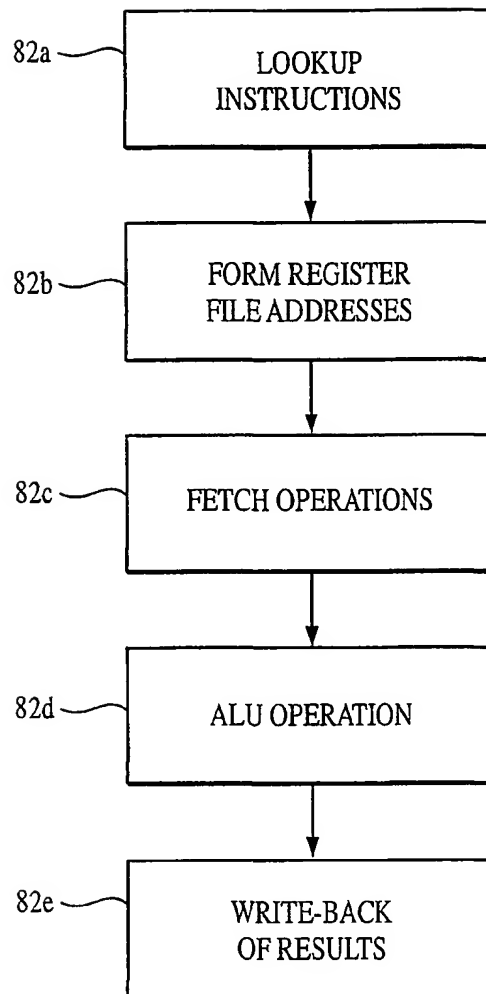
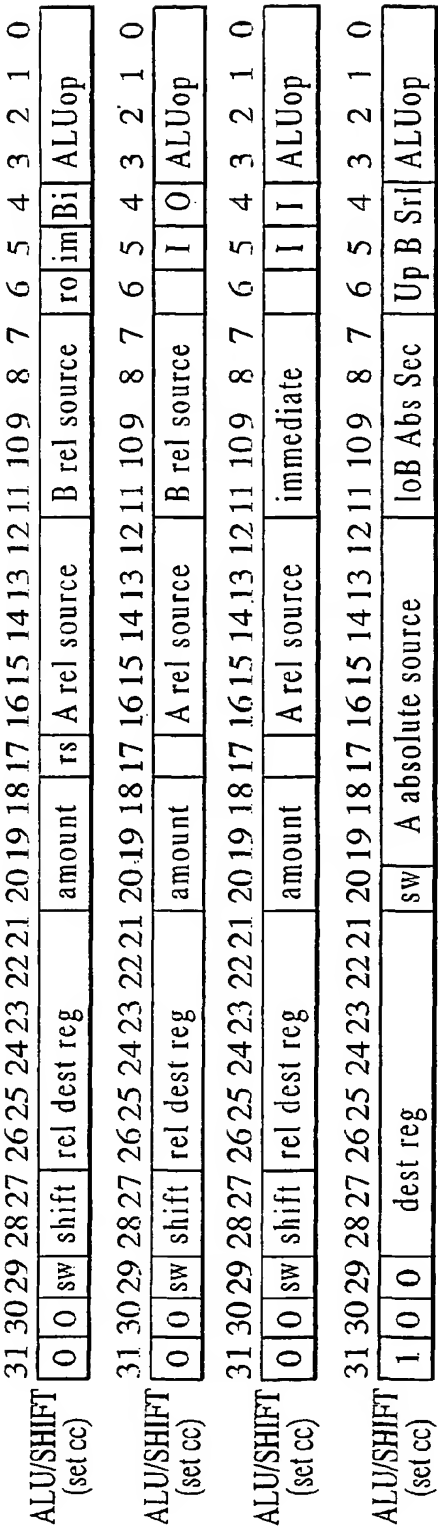


FIG. 4



Shift Decode:
(rs,r0) decode ([31:0] shifts into [63:32] and take [63:32]):
00 = left rotate
01 = right shift (32-ShfAmt = Right Shift Amt)
10 = left shift
11 = double shift (upper A-op shifts into lower B-op)
====> "left rotate" of zero gives zero shift (therwise zero amount signifies indirect shift)

ALU-OP decode:

0000 = B	0100 = ~A&B (~and)	1000 = A-B	1100 = A+B(8)
0001 = ~B	0101 = XOR	1001 = B-A	1101 = A+B(16)
0010 = A&B (and)	0110 = OR	1010 =	1110 = A+B
0011 = A&~B (and~)	0111 = mul-stuff	1011 =	1111 = A+B+Cin

FIG. 5

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US00/24000

A. CLASSIFICATION OF SUBJECT MATTER

IPC(7) : G06F 9/312, 9/315, 9/38
US CL : 712/43, 224, 225, 228, 229

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 712/43, 224, 225, 228, 229

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

USPAT, EPO, JPO, ACM, IEEE, Derwent, IBM Technical Disclosure Bulletins

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X,P	US 6,058,465 A (NGUYEN) 02 May 2000, col, 39-40 and 91-92.	1-2, 5-6, 8-10, 20-21, 24-25, and 27-29
A	US 6,002,881 A (YORK et al.) 14 December 1999	1-38
A	Waldspurger et al. Register Relocation: Flexible Contexts for Multithreading Proceedings of the 20th Annual International Symposium on Computer Architecture 1993 pages 120-130.	1-38



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents:	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"A" document defining the general state of the art which is not considered to be of particular relevance	"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"E" earlier document published on or after the international filing date	"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"&" document member of the same patent family
"O" document referring to an oral disclosure, use, exhibition or other means	
"P" document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search

14 DECEMBER 2000

Date of mailing of the international search report

08 JAN 2001

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer

WILLIAM M. KREAT

Telephone No. (703) 305-9699